

CISSP STUDY GUIDE

11th**HOUR**

- The only guide you need for last-minute studying
- Answers the toughest questions and highlights core topics
- Can be paired with any other study guide so you are completely prepared

Eric Conrad



Eleventh Hour CISSP®
Study Guide

Syngress Eleventh Hour Series

Eleventh Hour CISSP® Study Guide

ISBN: 978-1-59749-566-0

Eric Conrad

Eleventh Hour Linux+: Exam XK0-003 Study Guide

ISBN: 978-1-59749-497-7

Graham Speake

Eleventh Hour Security+: Exam SY0-201 Study Guide

ISBN: 978-1-59749-427-4

Ido Dubrawsky

Eleventh Hour Network+: N10-004 Study Guide

ISBN: 978-1-59749-428-1

Naomi Alpern

Visit www.syngress.com for more information on these titles and other resources.

Eleventh Hour CISSP® Study Guide

Lead Author

Eric Conrad

Contributing Authors

Seth Misener

Joshua Feldman

Technical Editor

Kevin Riggins



ELSEVIER

AMSTERDAM • BOSTON • HEIDELBERG • LONDON
NEW YORK • OXFORD • PARIS • SAN DIEGO
SAN FRANCISCO • SINGAPORE • SYDNEY • TOKYO

Syngress is an imprint of Elsevier

SYNGRESS®

Acquiring Editor: Rachel Roumeliotis
Development Editor: Matthew Cater
Project Manager: Marilyn E. Rash
Designer: Alisa Andreola

Syngress is an imprint of Elsevier
30 Corporate Drive, Suite 400
Burlington, MA 01803, USA

© 2011 Elsevier Inc. All rights reserved.

No part of this publication may be reproduced or transmitted in any form or by any means, electronic or mechanical, including photocopying, recording, or any information storage and retrieval system, without permission in writing from the publisher. Details on how to seek permission, further information about the Publisher's permissions policies and our arrangements with organizations such as the Copyright Clearance Center and the Copyright Licensing Agency, can be found at our website: www.elsevier.com/permissions.

This book and the individual contributions contained in it are protected under copyright by the Publisher (other than as may be noted herein).

Notices

Knowledge and best practice in this field are constantly changing. As new research and experience broaden our understanding, changes in research methods or professional practices may become necessary.

Practitioners and researchers must always rely on their own experience and knowledge in evaluating and using any information or methods described herein. In using such information or methods they should be mindful of their own safety and the safety of others, including parties for whom they have a professional responsibility.

To the fullest extent of the law, neither the Publisher nor the authors, contributors, or editors assume any liability for any injury and/or damage to persons or property as a matter of products liability, negligence or otherwise, or from any use or operation of any methods, products, instructions, or ideas contained in the material herein.

Library of Congress Cataloging-in-Publication Data

Conrad, Eric, 1967–

Eleventh hour CISSP® Study Guide / Eric Conrad.

p. cm.

ISBN 978-1-59749-566-0 (pbk.)

1. Computer networks—Security measures—Examinations—Study guides. 2. Computer security—Examinations—Study guides. 3. Telecommunications engineers—Certification.

I. Title.

TK5105.59.C579185 2011

005.8076–dc22

2010038528

British Library Cataloguing-in-Publication Data

A catalogue record for this book is available from the British Library.

For information on all Syngress publications
visit our website at www.syngress.com

Printed in the United States of America

10 11 12 13 14 10 9 8 7 6 5 4 3 2 1

Working together to grow
libraries in developing countries

www.elsevier.com | www.bookaid.org | www.sabre.org

ELSEVIER

BOOK AID
International

Sabre Foundation

Contents

ABOUT THE AUTHORS	vii
CHAPTER 1 Domain 1: Information Security Governance and Risk Management	1
CHAPTER 2 Domain 2: Access Control	19
CHAPTER 3 Domain 3: Cryptography	39
CHAPTER 4 Domain 4: Physical (Environmental) Security	55
CHAPTER 5 Domain 5: Security Architecture and Design	69
CHAPTER 6 Domain 6: Business Continuity and Disaster Recovery Planning	89
CHAPTER 7 Domain 7: Telecommunications and Network Security	109
CHAPTER 8 Domain 8: Application Development Security	129
CHAPTER 9 Domain 9: Operations Security	147
CHAPTER 10 Domain 10: Legal, Regulations, Investigations, and Compliance	161
GLOSSARY	175
INDEX	183

This page intentionally left blank

About the Authors

Eric Conrad, lead author (CISSP®, GIAC GSE, GPEN, GCIH, GCIA, GCFA, GAWN, GSEC, Security+), is a SANS-certified instructor and president of Backshore Communications, which provides information warfare, penetration testing, incident handling, and intrusion detection consulting services. Eric started his professional career in 1991 as a UNIX systems administrator for a small oceanographic communications company. He gained information security experience in a variety of industries, including research, education, power, Internet, and healthcare, in positions ranging from systems programmer to security engineer to HIPAA security officer and ISSO. He has taught more than a thousand students in courses such as SANS Management 414: CISSP®, Security 560: Network Penetration Testing and Ethical Hacking, Security 504: Hacker Techniques, and Exploits and Incident Handling.

Eric graduated from the SANS Technology Institute with a Master of Science degree in Information Security Engineering. He lives in Peaks Island, Maine, with his wife Melissa and children Eric and Emma.

Seth Misenaar, contributing author (CISSP®, GPEN, GCIH, GCIA, GCFA, GWAPT, GCWN, GSEC, MCSE, MCDBA), is a certified instructor with the SANS Institute and serves as lead consultant for Context Security, which is based in Jackson, Mississippi. His background includes security research, network and Web application penetration testing, vulnerability assessment, regulatory compliance, security architecture design, and general security consulting. Seth previously served as a physical and network security consultant for Fortune 100 companies and as the HIPAA and information security officer for a state government agency. He teaches a variety of courses for the SANS Institute, including Security Essentials, Web Application Penetration Testing, Hacker Techniques, and the CISSP® course.

Seth is pursuing a Master of Science degree in Information Security Engineering from the SANS Technology Institute and holds a Bachelor of Science degree from Millsaps College, Jackson, Mississippi. He resides in Jackson with his wife Rachel and children Jude and Hazel.

Joshua Feldman, contributing author (CISSP®), is currently employed by SAIC, Inc. He has been involved in the Department of Defense Information Systems Agency (DISA) Information Assurance Education, Training, and Awareness program since 2002, where he has contributed to a variety of DoD-wide Information Assurance and Cyber Security policies, specifically the 8500.2 and 8570 series. Joshua has taught more than a thousand DoD students through his “DoD IA Boot Camp” course. He is a subject matter expert for the Web-based DoD Information Assurance Awareness—yearly training of every DoD user is

required as part of his or her security awareness curriculum. Also, he is a regular presenter and panel member at the annual Information Assurance Symposium hosted jointly by DISA and NSA.

Before joining the support team at DoD/DISA, Joshua spent time as an IT security engineer at the Department of State's Bureau of Diplomatic Security. He got his start in the IT security field with NFR Security Software, a company that manufactures Intrusion Detection Systems. There, he worked as both a trainer and an engineer, implementing IDS technologies and instructing customers in properly configuring them.

Kevin Riggins, editor (CISSP®), has more than 22 years of experience in information technology and has focused on information security since 1999. He has been a certified information systems security professional since 2004 and currently works for a Fortune 500 financial service company, where he leads a team of information security analysts responsible for internal consulting, risk assessment, and vendor security review. Kevin writes about various information security topics on his blog, Infosec Ramblings (www.infosecramblings.com), has been published in *(IN)Secure* magazine, and is a frequent speaker at conference and industry association meetings.

CHAPTER 1

Domain 1: Information Security Governance and Risk Management

1

Exam Objectives in this Chapter

- Risk analysis
- Information security governance

INTRODUCTION

Our job as information security professionals is to evaluate risks against our critical assets and deploy safeguards to mitigate them. We work in various roles as firewall engineers, penetration testers, auditors, management, and the like. The common thread is risk: It is part of our job description.

The Information Security Governance and Risk Management domain focuses on risk analysis and mitigation. It also details security governance, or the organizational structure required for a successful information security program. The difference between organizations that are successful and those that fail in this realm is usually not tied to dollars or staff size.: It is tied to the right people in the right roles. Knowledgeable and experienced information security staff and supportive and vested leadership are the keys to success.

Speaking of leadership, learning to speak the language of leaders is another key to personal success in this industry. The ability to effectively communicate information security concepts with C-level executives is a rare and needed skill. This domain also helps you speak this language by discussing risk in terms such as Total Cost of Ownership (TCO) and Return on Investment (ROI).

RISK ANALYSIS

All information security professionals assess risk: We do it so often that it becomes second nature. A patch is released on a Tuesday. Your company normally tests for two weeks before installing, but a network-based worm is spreading on the Internet that infects unpatched systems. If you install the patch now, you risk downtime due to lack of testing. If you wait to test, you risk infection

by the worm. What is the bigger risk? What should you do? Risk Analysis (RA) will help you decide.

The average person does a poor job of accurately analyzing risk: If you fear the risk of dying while traveling and, to mitigate that risk, drive from New York to Florida instead of flying, you have done a poor job of analyzing risk. It is far riskier, per mile, to travel by car than by airplane when considering the risk of death while traveling.

Accurate Risk Analysis is a critical skill for an information security professional. We must hold ourselves to a higher standard when judging risk. Our risk decisions dictate which safeguards we deploy to protect our assets, and the amount of money and resources we spend doing so. Poor decisions result in wasted money or, even worse, compromised data.

Assets

Assets are the valuable resources you are trying to protect. They can be data, systems, people, buildings, property, and so forth. The value or criticality of the asset dictates the safeguards you deploy. People are your most valuable asset.

Threats and vulnerabilities

A **threat** is a potentially harmful occurrence, such as an earthquake, a power outage, or a network-based worm like Conficker (aka Downadup or Kido; see www.microsoft.com/security/worms/Conficker.aspx), which began attacking Microsoft Windows operating systems in late 2008. A threat is a negative action that may harm a system.

A **vulnerability** is a weakness that allows a threat to cause harm. Examples of vulnerabilities (matching our previous threats) are buildings that are not built to withstand earthquakes, a data center without proper backup power, or a Microsoft Windows XP system that has not been patched in a few years.

Fast Facts

Using the worm example, the threat is Conficker; it spreads through three vectors:

1. Lack of the MS08-067 patch (see www.microsoft.com/technet/security/Bulletin/MS08-067.mspx).
2. Infected USB tokens that “autorun” when inserted into a Windows system
3. Weak passwords on network shares.

A networked Microsoft Windows system is vulnerable if it lacks the patch, if it automatically runs software on a USB token when inserted, or if it has a network share with a weak password. If any of those three conditions are true, you have risk. A Linux system has no vulnerability to Conficker and therefore runs no risk from it.

Risk = Threat × Vulnerability

To have risk, a threat must connect to a vulnerability. This relationship is stated by the formula:

$$\text{Risk} = \text{Threat} \times \text{Vulnerability}$$

You can choose a value to specific risks using this formula. Assign a number to both threats and vulnerabilities. A common range is 1 through 5 (the range is arbitrary; just keep it consistent when comparing different risks).

Impact

The “Risk = Threat × Vulnerability” equation sometimes uses an added variable, **impact**: “Risk = Threat × Vulnerability × Impact.” Impact is the severity of the damage, sometimes expressed in dollars, which is why Risk = Threat × Vulnerability × Cost is sometimes used. A synonym for impact is consequences.

Exam Warning

Loss of human life has a near-infinite significance on the exam. When calculating risk using the “Risk = Threat × Vulnerability × Impact” formula, any risk involving loss of human life is extremely high and must be mitigated.

Risk Analysis Matrix

The Risk Analysis Matrix uses a quadrant to map the likelihood of a risk occurring against the consequences (or impact) that the risk would have. The Australia/New Zealand 4360 Standard on Risk Management (AS/NZS 4360, see *www.standards.org.au*) describes the Risk Analysis Matrix, which is shown in Table 1.1.

The Risk Analysis Matrix allows you to perform **Qualitative Risk Analysis** (see the section Qualitative and Quantitative Risk Analysis to come) based on likelihood (from rare to almost certain) and consequences, or impact,

		Consequences				
		Insignificant 1	Minor 2	Moderate 3	Major 4	Catastrophic 5
Likelihood	5. Almost certain	H	H	E	E	E
	4. Likely	M	H	H	E	E
	3. Possible	L	M	H	E	E
	2. Unlikely	L	L	M	H	E
	1. Rare	L	L	M	H	H

(from insignificant to catastrophic). The resulting risk scores are Low (L), Medium (M), High (H), and Extreme (E). Low risks are handled via normal processes; moderate risks require management notification; high risks require senior management notification; and extreme risks require immediate action, including a detailed mitigation plan (and senior management notification).

The goal of the matrix is to identify high-likelihood/high-consequence risks (upper right quadrant of Table 1.1) and drive them down to the low-likelihood/low-consequence level (lower left quadrant).

Calculating Annualized Loss Expectancy

The Annualized Loss Expectancy (ALE) calculation allows you to determine the annual cost of a loss due to a given risk. Once calculated, ALE allows you to make informed decisions to mitigate the risk.

This section uses an example of risk due to lost or stolen unencrypted laptops. Assume that your company has 1000 laptops that contain Personally Identifiable Information (PII). You are the Security Officer, and your concern is the risk of exposure of PII due to the laptops' misplacement or theft. You want to purchase and deploy a laptop encryption solution. The solution is expensive, so you need to convince management that it is worthwhile.

ASSET VALUE

The Asset Value (AV) is the value of the asset you are trying to protect. In this example, each laptop costs \$2,500, but the real value is in the PII it contains. Theft of unencrypted PII occurred previously and cost the company many times the value of the laptops in regulatory fines, bad publicity, legal fees, staff hours spent investigating, and so forth. The true average Asset Value of a laptop with PII for this example is \$25,000 (\$2,500 for the hardware and \$22,500 for the exposed PII).

EXPOSURE FACTOR

The Exposure Factor (EF) is the percentage of value lost by an asset because of an incident. In the case of a stolen laptop with unencrypted PII, the Exposure Factor is 100%: The laptop and all the data are gone.

SINGLE LOSS EXPECTANCY

The Single Loss Expectancy (SLE) is the cost of a single loss. SLE is the Asset Value (AV) times the Exposure Factor (EF). In our case, SLE is \$25,000 (Asset Value) times 100% (Exposure Factor), or \$25,000.

ANNUAL RATE OF OCCURRENCE

The Annual Rate of Occurrence (ARO) is the number of losses you suffer per year. Looking through past events, you discover that you have suffered 11 lost or stolen laptops per year on average, so your ARO is 11.

Table 1.2 Summary of Risk Equations

	Formula	Description
Asset Value (AV)	AV	Value of asset
Exposure Factor (EF)	EF	Percentage of asset value lost
Single Loss Expectancy (SLE)	$AV \times EF$	Cost of one loss
Annual Rate of Occurrence (ARO)	ARO	Number of losses per year
Annualized Loss Expectancy (ALE)	$SLE \times ARO$	Cost of losses per year

ANNUALIZED LOSS EXPECTANCY

The Annualized Loss Expectancy (ALE) is your yearly cost due to a risk. It is calculated by multiplying the Single Loss Expectancy (SLE) times the Annual Rate of Occurrence (ARO). In our case it is \$25,000 (SLE) times 11 (ARO), or \$275,000.

Table 1.2 summarizes the equations used to determine Annualized Loss Expectancy.

Total Cost of Ownership

The Total Cost of Ownership (TCO) is the total cost of a mitigating safeguard. It combines upfront costs (often one-time capital expenses) and annual cost of maintenance, including staff hours, vendor maintenance fees, software subscriptions, and so forth. These ongoing costs are usually considered operational expenses.

Using our laptop encryption example, the upfront cost of laptop encryption software is \$100/laptop, or \$100,000 for 1,000 laptops. The vendor charges a 10% annual support fee, or \$10,000/year. You estimate that it will take 4 staff hours per laptop to install the software, or 4,000 staff hours in total. The staff that performs this work makes \$50/hour plus benefits. Including benefits, the staff cost per hour is \$70 times 4,000 hours, or \$280,000.

Your company uses a three-year technology refresh cycle, so you calculate the Total Cost of Ownership over three years:

- Software cost: \$100,000
- Three years of vendor support: $\$10,000 \times 3 = \$30,000$
- Hourly staff cost: \$280,000
- Total Cost of Ownership over three years: \$410,000
- Total Cost of Ownership per year: $\$410,000/3 = \$136,667/\text{year}$

Your Annual Total Cost of Ownership for the laptop encryption project is \$136,667 per year.

Return on Investment

The Return on Investment (ROI) is the amount of money saved by implementing a safeguard. If your annual Total Cost of Ownership (TCO) is less than your

Table 1.3 Annualized Loss Expectancy for Unencrypted Laptops

	Formula	Value
Asset Value (AV)	AV	\$25,000
Exposure Factor (EF)	EF	100%
Single Loss Expectancy (SLE)	$AV \times EF$	\$25,000
Annual Rate of Occurrence (ARO)	ARO	11
Annualized Loss Expectancy (ALE)	$SLE \times ARO$	\$275,000

Table 1.4 Annualized Loss Expectancy for Encrypted Laptops

	Formula	Value
Asset Value (AV)	AV	\$25,000
Exposure Factor (EF)	EF	10%
Single Loss Expectancy (SLE)	$AV \times EF$	\$2,500
Annual Rate of Occurrence (ARO)	ARO	11
Annualized Loss Expectancy (ALE)	$SLE \times ARO$	\$27,500

Annualized Loss Expectancy (ALE), you have a positive ROI (and have made a good choice). If your TCO is higher than your ALE, you have made a poor choice.

The annual TCO of laptop encryption is \$136,667; the Annualized Loss Expectancy for lost or stolen unencrypted laptops is \$275,000. The math is summarized in Table 1.3.

Implementing laptop encryption will change the Exposure Factor. The laptop hardware is worth \$2,500, and the exposed PII costs an additional \$22,500, for a \$25,000 Asset Value. If an unencrypted laptop is lost or stolen, the EF is 100% (the hardware and all data are exposed). Laptop encryption mitigates the PII exposure risk, lowering the exposure factor from 100% (the laptop and all data) to 10% (just the laptop hardware).

The lower Exposure Factor lowers the Annualized Loss Expectancy from \$275,000 to \$27,500, as shown in Table 1.4.

You will save \$247,500/year (the old ALE, \$275,000, minus the new ALE, \$27,500) by making an investment of \$136,667. Your ROI is \$110,833 per year (\$247,500 minus \$136,667). The laptop encryption project has a positive ROI and is a wise investment.

Risk choices

Once we have assessed risk, we must decide what to do. Options include accepting the risk, mitigating or eliminating it, transferring it, and avoiding it.

ACCEPT THE RISK

Some risks may be accepted: In certain cases, it is cheaper to leave an asset unprotected from a specific risk rather than make the effort (and spend the money) required to protect it. This cannot be an ignorant decision: The risk, and all options, must be considered before you can accept it.

Risk Acceptance Criteria

Low-likelihood/low-consequence risks are candidates for risk acceptance. High and Extreme risks are not. There are cases, such as data protected by laws or regulations or risk to human life or safety, where accepting the risk is not an option.

MITIGATE THE RISK

Mitigating the risk means lowering it to an acceptable level. The laptop encryption example given previously in the Annualized Loss Expectancy section is an example of risk mitigation. The risk of lost PII due to stolen laptops was mitigated by encrypting the data on them. It was not eliminated entirely: A weak or exposed encryption password could expose the PII, but the risk was reduced to an acceptable level.

In some cases it is possible to remove the risk entirely: this is called eliminating it.

TRANSFER THE RISK

Risk transfer is the "insurance model." Most people do not assume the risk of fire to their house: They pay an insurance company to assume that risk for them.

AVOID THE RISK

A thorough Risk Analysis should be completed before taking on a new project. If it discovers high or extreme risks that cannot be easily mitigated, avoiding the risk (and the project) may be the best option.

The math for this decision is straightforward: Calculate the Annualized Loss Expectancy of the new project, and compare it with the Return on Investment expected from the project. If the ALE is higher than the ROI (even after risk mitigation), risk avoidance is the best course. There may also be legal or regulatory reasons that will dictate avoiding the risk.

Qualitative and Quantitative Risk Analysis

Quantitative and **Qualitative** Risk Analysis are two methods for analyzing risk. Quantitative Risk Analysis uses hard metrics, such as dollars. Qualitative Risk Analysis uses simple approximate values. Quantitative is more objective; qualitative is more subjective.

Exam Warning

Quantitative Risk Analysis requires you to calculate the quantity of the asset you are protecting. "Quantitative-quantity" is a hint to remember on the exam.

Calculating the Annualized Loss Expectancy (ALE) is an example of Quantitative Risk Analysis. The inputs for ALE are hard numbers: Asset Value (in dollars), Exposure Factor (as a percentage), and Annual Rate of Occurrence (as a hard number).

The Risk Analysis Matrix (shown in Table 1.1) is an example of Qualitative Risk Analysis. Likelihood and Consequences are rough (and sometimes subjective) values, ranging from 1 to 5. Whether the consequences of a certain risk are a 4 or a 5 can be a matter of (subjective) debate.

Quantitative Risk Analysis is more difficult: To quantitatively analyze the risk of damage to a data center due to an earthquake, you need to calculate the asset value of the data center, such as the cost of the building, servers, network equipment, computer racks, and monitors. Then you must calculate the Exposure Factor and so on.

To qualitatively analyze the same risk, you research it and agree that the likelihood is a 2 and the consequences are a 4, and use the Risk Analysis Matrix to determine a risk of “high.”

The risk management process

Published by the National Institute of Standards and Technology (NIST), Special Publication 800-30, Risk Management Guide for Information Technology Systems (see: <http://csrc.nist.gov/publications/nistpubs/800-30/sp800-30.pdf>), describes a nine-step Risk Analysis process:

- 1.** System Characterization
- 2.** Threat Identification
- 3.** Vulnerability Identification
- 4.** Control Analysis
- 5.** Likelihood Determination
- 6.** Impact Analysis
- 7.** Risk Determination
- 8.** Control Recommendations
- 9.** Results Documentation¹

We covered these steps individually; let’s end this section by following NIST’s process.

Step 1, System Characterization, describes the scope of the risk management effort and the systems that will be analyzed. Threat Identification and Vulnerability Identification, Steps 2 and 3, identify the threats and vulnerabilities required to determine risks using the “Risk = Threat × Vulnerability” formula.

Step 4, Control Analysis, analyzes the security controls (safeguards) in place or planned to mitigate risk. Steps 5 and 6, Likelihood Determination and Impact Analysis, identify important risks (especially those with high likelihood and high impact/consequence).

Steps 1 through 7 are used to determine Control Recommendations, or the risk mitigation strategy. That strategy is documented in Step 8, Results Documentation.

INFORMATION SECURITY GOVERNANCE

Information Security Governance is information security at the organizational level: senior management, policies, processes, and staffing. It is also the organizational priority, provided by senior leadership, that is required for a successful information security program.

Security policy and related documents

Documents such as policies and procedures are a required part of any successful information security program. These should be grounded in reality: They are not idealistic documents that sit on shelves collecting dust, but should mirror the real world and provide guidance on the correct (and sometimes required) way of doing things.

POLICY

Policies are high-level management directives. Policy is mandatory: If you do not agree with your company's policy on sexual harassment, for example, you do not have the option of not following it.

Crunch Time

Policy is high level: It does not delve into specifics. A server security policy discusses protecting the confidentiality, integrity, and availability of the system (usually in those terms). It may discuss software updates and patching. It

does not use terms like "Linux" or "Windows"; these are too low level. In fact, if you converted your servers from Windows to Linux, your server policy would not change. Other documents, like procedures, would change.

Components of Program Policy

All policy should contain these basic components:

- Purpose
- Scope
- Responsibilities
- Compliance

Purpose describes the need for the policy, typically to protect the confidentiality, integrity, and availability of data.

Scope describes the systems, people, facilities, and organizations that are covered by the policy. Any related entities that are not in scope should be documented to avoid confusion.

Responsibilities include those of the information security staff and policy and management teams, as well as those of all members of the organization.

Compliance describes two related issues: how to judge the effectiveness of the policies (how well they are working), and what happens when a policy is violated (the sanction). All policy must have “teeth”: One that forbids accessing explicit content via the Internet is not useful if there are no consequences for doing so.

Policy Types

The NIST Special Publication 800-12 (see <http://csrc.nist.gov/publications/nistpubs/800-12/800-12-html/chapter5.html>) discusses three specific policy types: program, issue-specific, and system-specific.

Program policy establishes an organization’s information security program. Examples of issue-specific policies listed in NIST SP 800-12 include email and email privacy. System-specific policy examples include file servers and web servers.

PROCEDURES

A **procedure** is a step-by-step guide for accomplishing a task. Procedures are low-level and specific. Like policies, they are mandatory.

Here is a simple example of a procedure for creating a new user:

1. Receive a new-user request form and verify its completeness.
2. Verify that the user’s manager has signed the form.
3. Verify that the user has read and agreed to the user account security policy.
4. Classify the user’s role by following role-assignment procedure NX-103.
5. Verify that the user has selected a “secret word,” such as mother’s maiden name, and enter it into the help desk account profile.
6. Create the account and assign the proper role.
7. Assign the secret word as the initial password, and set “Force user to change password on next login” to “True.”
8. Email the New-Account document to the user and his or her manager.

The steps in this procedure are mandatory. Security administrators do not have the option of skipping Step 1, for example, and creating an account without a form.

DID YOU KNOW?

Other safeguards depend on the fact that procedures are mandatory: When a user calls the help desk because of a forgotten password, the help desk follows its “forgotten password” procedure, which includes asking for the user’s secret word. It cannot do that unless Step 5 was completed: Without that word, the help desk cannot securely reset the password. This mitigates social engineering attacks, in which an imposter tries to trick the help desk into resetting a password for an account he or she is not authorized to access.

STANDARDS

A **standard** describes the specific use of technology, often applied to hardware and software. “All employees will receive an ACME Nexus-6 laptop with 2 gigabytes of memory, a 2.8-GHZ duo core CPU, and 300-gigabyte disk” is an example of a hardware standard. “The laptops will run Windows 7 Professional, 32-bit version” is an example of a software (operating system) standard.

Standards are mandatory. They lower the Total Cost of Ownership of a safeguard. They also support disaster recovery, as recovering from disasters is easier when standards are employed.

GUIDELINES

Guidelines are recommendations (which are discretionary). They can include a useful piece of advice, such as “To create a strong password, take the first letter of every word in a sentence, and mix in some numbers and symbols. ‘I will pass the CISSP exam in 6 months!’ becomes ‘Iwptcei6m!’”

You can create a strong password without following this advice, which is why guidelines are not mandatory. They are useful, especially for novice users.

BASELINES

Baselines are uniform ways of implementing a safeguard. An example is “Harden the system by applying the Center for Internet Security Linux benchmarks” (see www.cisecurity.org for the CISecurity benchmarks; they are a great resource). The system must meet the baseline described by a benchmark.

Baselines are discretionary. It is acceptable to harden the system without following the aforementioned benchmarks, as long as it is at least as secure as a system hardened using the benchmarks.

Table 1.5 summarizes the types of security documentation.

Table 1.5 Summary of Security Documentation

Document	Example	Mandatory or Discretionary?
Policy	Protect the CIA of PII by hardening the operating system	Mandatory
Procedure	Step 1: Install prehardened OS Image. Step 2: Download patches from update server. Step 3: . . .	Mandatory
Standard	Use Nexus-6 laptop hardware	Mandatory
Guideline	Patch installation may be automated via the use of an installer script	Discretionary
Baselines	Use the CISecurity Windows Hardening benchmark	Discretionary

Security awareness and training

Security awareness and training are often confused. Awareness changes user behavior; training provides a skill set.

Reminding users to never share accounts or write their passwords down is an example of awareness. It is assumed that some users are doing the wrong thing, and awareness is designed to change that behavior.

Security training teaches a user how to do something. Examples include training new help desk personnel to open, modify, and close service tickets; training network engineers to configure a router; and training a security administrator to create a new account.

Roles and responsibilities

Primary information security roles include Senior Management, Data Owner, Custodian, and User. Each plays a different role in securing an organization's assets.

Senior Management creates the information security program and ensures that it is properly staffed and funded, and has organizational priority. It is responsible for ensuring that all organizational assets are protected.

The **Data Owner** (also called the Information Owner or the Business Owner) is a management employee responsible for ensuring that specific data is protected. Data Owners determine data sensitivity labels and how frequently data should be backed up. A company with multiple lines of business may have multiple Data Owners. The Data Owner performs management duties, whereas the hands-on protection of data is performed by Custodians.

A Custodian provides hands-on protection of assets such as data. He or she performs data backups and restoration, patches systems, configures antivirus software, and so forth. Custodians follow detailed orders; they do not make critical decisions on how data is protected. The Data Owner may dictate that "All data must be backed up every 24 hours." The Custodians (and their managers) then deploy and operate a backup solution that meets the Data Owner's requirements.

The **User** is the fourth primary information security role. Users must follow the rules: They must comply with mandatory policies, procedures, standards, and so forth. They must not write their passwords down or share accounts, for example. Users must be made aware of these risks and requirements. You cannot assume they will know what to do or that they are already doing the right thing: They must be told, via information security awareness.

Privacy

Privacy is the protection of the confidentiality of personal information. Many organizations host personal information about their users: PII such as social security numbers, financial information (e.g., annual salary and bank account

information required for payroll deposits), and healthcare information for insurance purposes. The confidentiality of this information must be assured.

Outsourcing and offshoring

Outsourcing is using a third party to provide Information Technology support services that were previously performed in-house. **Offshoring** is outsourcing to another country.

Both can lower Total Cost of Ownership by providing IT services at lower cost. They may also enhance the information technology resources and skill set and resources available to a company (especially a small one), which can improve confidentiality, integrity, and availability of data.

Offshoring can raise privacy and regulatory issues. For a U.S. company with data offshored to Australia, there is no Health Insurance Portability and Accountability Act (HIPAA), the primary regulation covering healthcare data in the United States. Nor is there SOX (Sarbanes-Oxley), for example, which protects publicly traded data in the United States, or GLBA (Gramm-Leach-Bliley Act), which protects financial information in the United States), and so forth.

A thorough and accurate Risk Analysis must be performed before outsourcing or offshoring sensitive data. If the data will reside in another country, you must ensure that laws and regulations governing it are followed, even beyond the country's jurisdiction. This can be done contractually: The Australian outsourcing company can, for example, agree to follow HIPAA via contract.

Auditing and control frameworks

Auditing means verifying compliance to a security control framework (or published specification). It helps support Risk Analysis efforts by verifying that a company not only “talks the talk” (has documentation supporting a robust information security program) but also “walks the walk” (actually has a robust information security program in practice).

A number of control frameworks are available to assist the auditing of Risk Analysis. Some, such as PCI (the Payment Card Industry), are industry-specific (vendors who use credit cards in the example). Others, such as OCTAVE, ISO 17799/27002, and COBIT, covered next, are more general.

OCTAVE

OCTAVE is the Operationally Critical Threat, Asset, and Vulnerability Evaluation, a risk management framework from Carnegie Mellon University. It describes a three-phase process for managing risk. Phase 1 identifies staff knowledge, assets, and threats. Phase 2 identifies vulnerabilities and evaluates safeguards. Phase 3 conducts the Risk Analysis and develops the risk mitigation strategy.

OCTAVE is a high-quality free resource that may be downloaded from www.cert.org/octave/.

ISO 17799 AND THE ISO 27000 SERIES

ISO 17799 was a broad-based approach for the information security code of practice of the International Organization for Standardization (based in Geneva, Switzerland). The full title is “ISO/IEC 17799:2005 Information Technology: Security Techniques—Code of Practice for Information Security Management.” “ISO 17799:2005” signifies the 2005 version of the standard. It was based on BS (British Standard) 7799 part one.

Fast Facts

ISO 17799 had 11 areas, focusing on specific information security controls:

1. Policy
2. Organization of information security
3. Asset management
4. Human resources security
5. Physical and environmental security
6. Communications and operations management
7. Access control
8. Information systems acquisition, development, and maintenance
9. Information security incident management
10. Business continuity management
11. Compliance²

ISO 17799 was renumbered as ISO 27002 in 2005 to make it consistent with the 27000 series of ISO security standards. ISO 27001 is a related standard, formally called ISO/IEC 27001:2005 Information technology: Security techniques—Information Security Management Systems—Requirements. It was based on BS 7799 part two.

Note that the title of ISO 27002 includes the word “Techniques”; ISO 27001 includes the word “Requirements.” Simply put, ISO 27002 describes information security best practices (techniques), and ISO 27001 describes a process for auditing those best practices (requirements).

COBIT

COBIT (Control Objectives for Information and Related Technology) is a control framework for employing information security governance best practices within an organization. COBIT’s purpose is

to provide management and business process owners with an information technology (IT) governance model that helps in delivering value from IT and understanding and managing the risks associated with IT. COBIT helps bridge the gaps amongst business requirements, control needs and technical issues. It is a control model to meet the needs of IT governance and ensure the integrity of information and information systems.³

COBIT has 34 IT processes. More information about it is available at www.isaca.org/cobit/. COBIT was developed by ISACA (Information Systems Audit and Control Association; see www.isaca.org). Version 4.1 was released in 2007.

ITIL

ITIL (Information Technology Infrastructure Library) is a framework for providing best services in IT Service Management (ITSM). More information about it is available at www.itil-officialsite.com.

ITIL contains five “Service Management Practices–Core Guidance” publications:

- Service strategy
- Service design
- Service transition
- Service operation
- Continual service improvement

Service Strategy helps IT provide services. Service Design details the infrastructure and architecture required to deliver IT services. Service Transition describes taking new projects and making them operational. Service Operation covers IT operations controls. Finally, Continual Service Improvement describes ways to improve existing IT services.

Certification and accreditation

Certification is a detailed inspection that verifies whether a system meets the documented security requirements. **Accreditation** is the Data Owner’s acceptance of the risk represented by a system. This process is called Certification and Accreditation, or C&A.

The NIST Special Publication 800-37, “Guide for the Security Certification and Accreditation of Federal Information Systems” (see <http://csrc.nist.gov/publications/nistpubs/800-37-rev1/sp800-37-rev1-final.pdf>), describes U.S. federal certification and accreditation.

According to NIST,

Security certification is a comprehensive assessment of the management, operational, and technical security controls in an information system, made in support of security accreditation, to determine the extent to which the controls are implemented correctly, operating as intended, and producing the desired outcome with respect to meeting the security requirements for the system.

Also,

Security accreditation is the official management decision given by a senior agency official to authorize operation of an information system and to explicitly accept the risk to agency operations, agency assets, or individuals based on the implementation of an agreed-upon set of security controls.⁴

Certification may be performed by a trusted third party such as an auditor. Certifiers investigate a system, inspect documentation, and may observe operations. They audit the system to ensure compliance. Certification is only a recommendation: The certifier does not have the ability to approve a system or an environment. Only the Data Owner (the Accreditor) can do so.

NIST SP 800-37 describes a four-step Certification and Accreditation process:

- Initiation Phase
- Security Certification Phase
- Security Accreditation Phase
- Continuous Monitoring Phase

The information security system and risk mitigation plan is researched during the initiation phase. The security of the system is assessed and documented during the security certification phase. The decision to accept the risk represented by the system is made and documented during the security accreditation phase. Finally, once accredited, the ongoing security of the system is verified during the continuous monitoring phase.

SUMMARY OF EXAM OBJECTIVES

Information security governance ensures that an organization has the correct information structure, leadership, and guidance. Risk Analysis (RA) helps ensure that an organization properly identifies, analyzes, and mitigates risk. All three of these qualities—information security governance, ethics, and Risk Analysis—are crucial for the success of an organization.

Finally, accurately assessing risk and understanding terms such as Annualized Loss Expectancy, Total Cost of Ownership, and Return on Investment will not only help you on the exam but also help advance your information security career.

TOP FIVE TOUGHEST QUESTIONS

1. Which control framework has 34 IT processes?
 - A. COSO
 - B. COBIT
 - C. ITIL
 - D. OCTAVE
2. What is the difference between a standard and a guideline?
 - A. Standards are compulsory; guidelines are mandatory.
 - B. Standards are recommendations; guidelines are requirements.
 - C. Standards are requirements; guidelines are recommendations.
 - D. Standards are recommendations; guidelines are optional.
3. Which phase of OCTAVE identifies vulnerabilities and evaluates safeguards?
 - A. Phase 1
 - B. Phase 2

- C. Phase 3
 - D. Phase 4
4. What was ISO 17799 renamed as?
- A. BS 7799-1
 - B. ISO 27000
 - C. ISO 27001
 - D. ISO 27002
5. Which of the following describes a duty of the Data Owner?
- A. Patch systems
 - B. Report suspicious activity
 - C. Ensure that files are backed up
 - D. Ensure that data has proper security labels

Answers

1. Correct Answer and Explanation: **B**. Answer **B** is correct; COBIT has 34 Information Technology processes.
Incorrect Answers and Explanations: **A**, **C**, and **D**. Answers **A**, **C**, and **D** are incorrect. All are audit or control frameworks, but only COBIT has 34 processes.
2. Correct Answer and Explanation: **C**. Answer **C** is correct; standards are requirements (mandatory); and guidelines are recommendations.
Incorrect Answers and Explanations: **A**, **B**, and **D**. Answers **A**, **B**, and **D** are incorrect. For **A**, guidelines are recommendations (*compulsory* and *mandatory* are synonyms). **B** has the recommendations and requirements flipped. For **D**, standards, not recommendations, are mandatory.
3. Correct Answer and Explanation: **B**. Answer **B** is correct; phase 2 identifies vulnerabilities and evaluates safeguards.
Incorrect Answers and Explanations: **A**, **C**, and **D**. Answers **A**, **C**, and **D** are incorrect. Phase 1 identifies staff knowledge, assets, and threats. Phase 3 conducts the Risk Analysis and develops the risk mitigation strategy. There is no Phase 4 in OCTAVE.
4. Correct Answer and Explanation: **D**. Answer **D** is correct; ISO 17799 was renamed as ISO 27002.
Incorrect Answers and Explanations: **A**, **B**, and **C**. Answers **A**, **B**, and **C** are incorrect. BS 7799-1 was the precursor to ISO 17799. ISO 27000 is a series of information security standards documents. ISO 270021 is another ISO 27000-series document designed to support auditing.
5. Correct Answer and Explanation: **D**. Answer **D** is correct; the Data Owner ensures that data has proper security labels.
Incorrect Answers and Explanations: **A**, **B**, and **C**. Answers **A**, **B**, and **C** are incorrect. Custodians patch systems. Users should be aware and report suspicious activity. Ensuring that files are backed up is a weaker answer for a Data Owner duty, used to confuse the Data Owner with the “owner of the file” on a discretionary access control system.

Endnotes

1. <http://csrc.nist.gov/publications/nistpubs/800-30/sp800-30.pdf> (accessed July 21, 2010).
2. www.iso.org/iso/catalogue_detail?csnumber=39612 (accessed July 21, 2010).
3. www.isaca.org/Knowledge-Center/cobit/Pages/FAQ.aspx (accessed July 21, 2010).
4. <http://csrc.nist.gov/publications/nistpubs/800-37-rev1/sp800-37-rev1-final.pdf> (accessed July 21, 2010).

CHAPTER 2

Domain 2: Access Control

Exam Objectives in this Chapter

- Cornerstone access control concepts
- Access control models
- Procedural issues for access control
- Access control defensive categories and types
- Authentication methods
- Access control technologies
- Assessing access control

INTRODUCTION

Access controls protect against threats such as unauthorized access, inappropriate modification of data, and loss of confidentiality. Access control is achieved by implementing strong technical, physical, and administrative measures.

Cornerstone access control concepts

Understanding cornerstone access control concepts, including confidentiality, integrity, and availability (as well as their mirror opposites: disclosure, alteration, and destruction), and subjects and objects, is a critical foundation for understanding access control.

Confidentiality, integrity, and availability

Confidentiality, Integrity, and Availability are the “CIA triad,” the cornerstone concept of information security. The triad, shown in Figure 2.1, forms the three-legged stool that information security is built on.

CONFIDENTIALITY

Confidentiality seeks to prevent the unauthorized disclosure of information. In other words, it seeks to prevent unauthorized read access to data. An example of a confidentiality attack would be the theft of **Personally Identifiable Information** (PII) such as credit card data.




FIGURE 2.1
The CIA triad.

INTEGRITY

Integrity seeks to prevent unauthorized modification of information. In other words, it seeks to prevent unauthorized write access to data.

AVAILABILITY

Availability ensures that information is at hand when needed. Systems need to be usable (available) for normal business use. An example of an attack on availability would be Denial of Service (DoS), which seeks to deny system service (or availability).

Crunch Time

There are two types of integrity: data and system. Data integrity seeks to protect information against unauthorized modification; system integrity seeks to

protect a system, such as a Windows 2008 server operating system, from unauthorized modification.

DISCLOSURE, ALTERATION, AND DESTRUCTION

The CIA triad may also be described by its opposite: **Disclosure**, **Alteration**, and **Destruction** (DAD). Disclosure is the unauthorized disclosure of information; alteration is the unauthorized modification of data; destruction is making systems unavailable. While the CIA acronym sometimes changes, the DAD acronym is shown in that order.

Identity and authentication, authorization, and accountability

The term “AAA” is often used to describe the cornerstone concepts **Authentication**, **Authorization**, and **Accountability**. Left out of the AAA acronym is **Identification**, which is required before the three As can follow.

IDENTITY AND AUTHENTICATION

Identity is a claim: If your name is Person X, you identify yourself by saying “I am Person X.” Identity alone is weak because there is no proof. You can also identify yourself by saying “I am Person Y.” Proving an identity claim is called authentication: You authenticate the identity claim, usually by supplying a piece of information or an object that only you possess, such as a password or your passport.

AUTHORIZATION

Authorization describes the actions you can perform on a system once you have been identified and authenticated. Actions may include reading, writing, and executing files or programs.

ACCOUNTABILITY

Accountability holds users accountable for their actions. This is typically done by logging and analyzing audit data. Enforcing accountability helps keep “honest people honest.” For some users, knowing that data is logged is not enough to provide accountability: They must know that the data is logged and audited, and that **sanctions** may result from violation of **policy**.

Nonrepudiation

Nonrepudiation means that a user cannot deny (repudiate) having performed a transaction. It combines authentication and integrity: Nonrepudiation authenticates the identity of a user who performs a transaction, and ensures the integrity of that transaction. You must have both authentication and integrity to have nonrepudiation: Proving you signed a contract to buy a car (authenticating your identity as the purchaser) is not useful if the car dealer can change the price from \$20,000 to \$40,000 (violating the integrity of the contract).

Least privilege and need to know

Least privilege means that users should be granted the minimum amount of access (authorization) required to do their jobs, and no more. **Need to know** is more granular than least privilege: The user must need to know that specific piece of information before accessing it.

Defense in depth

Defense in Depth (also called *layered defenses*) applies multiple safeguards (also called *controls*—that is, measures taken to reduce risk) to protect an asset. Any one security control may fail; by deploying multiple controls, you improve the confidentiality, integrity, and availability of your data.

Subjects and objects

A **subject** is an active entity on a data system. Most examples involve people accessing data files. However, running computer programs are subjects as well. A Dynamic Link Library file or a Perl script that updates database files with new information is also a subject.

An **object** is any passive data within the system. Objects can range from databases to text files. The important thing to remember about objects is that they are passive within the system. They do not manipulate other objects.

ACCESS CONTROL MODELS

Now that we have reviewed the cornerstone access control concepts, we can discuss the different access control models: The primary models are Discretionary Access Control (DAC), Mandatory Access Control (MAC), and Nondiscretionary Access Control.

Discretionary access control

Discretionary Access Control gives subjects full control of objects they have or have been given access to, including sharing the objects with other subjects. Subjects are empowered and control their data. Standard UNIX and Windows operating systems use DAC for filesystems: Subjects can grant other subjects access to their files, change their attributes, alter them, or delete them.

Mandatory access control

Mandatory Access Control is system-enforced, based on a subject's clearance and an object's labels. Subjects and objects have clearances and labels, respectively, such as confidential, secret, and top secret. A subject may access an object only if its clearance is equal to or greater than the object's label. Subjects cannot share objects with other subjects who lack the proper clearance, or "write down" objects to a lower classification level (e.g., from top secret to secret). MAC systems are usually focused on preserving the confidentiality of data.

Nondiscretionary access control

Role-Based Access Control (RBAC) defines how information is accessed on a system based on the role of the subject. A role can be a nurse, a backup administrator, a help desk technician, and the like. Subjects are grouped into roles, and each defined role has access permissions based on it, not on the individual.

RBAC is a type of nondiscretionary access control because users do not have discretion regarding the groups of objects they are allowed to access, and they are unable to transfer objects to other subjects.

Task-based access control is another nondiscretionary access control model related to RBAC. It is based on the tasks each subject must perform, such as writing prescriptions, restoring data from a backup tape, or opening a help desk ticket.

Centralized access control

Centralized access control is concentrated at one logical point for a system or organization. Instead of using local access control databases, systems authenticate via third-party authentication servers. Centralized access control can be used to provide Single Sign-On (SSO), where a subject may authenticate once and then access multiple systems.

Decentralized access control

Decentralized access control allows the IT administration to be closer to the mission and operations of the organization. With it, an organization spans multiple locations, and the local sites support and maintain independent systems, access control databases, and data. Decentralized access control is also called *distributed access control*.

Access control protocols and frameworks

Both centralized and decentralized models may support remote users authenticating to local systems. A number of protocols and frameworks may be used to support this need, including RADIUS, Diameter, TACACS/TACACS+, PAP, and CHAP.

RADIUS

The Remote Authentication Dial In User Service (**RADIUS**) protocol is a third-party authentication system. It is described in RFCs 2865 and 2866, and uses User Datagram Protocol (UDP) ports 1812 (authentication) and 1813 (accounting). RADIUS formerly used the (unofficially assigned) ports 1645 and 1646 for the same respective purposes; some continue to use them.

RADIUS is considered an “AAA” system, comprising three components: authentication, authorization, and accounting. It authenticates a subject’s credentials against an authentication database. It authorizes users by allowing specific users access to specific data objects. It accounts for each data session by creating a log entry for each RADIUS connection made. RADIUS request and response data is carried in Attribute Value Pairs (AVPs).

DIAMETER

Diameter is RADIUS’s successor, designed to provide an improved Authentication, Authorization, and Accounting (AAA) framework. RADIUS provides limited accountability and has problems with flexibility, scalability, reliability, and security. Diameter also uses Attribute Value Pairs, but supports many more: While RADIUS uses 8 bits for the AVP field (allowing 256 total possible AVPs), Diameter uses 32 bits for the AVP field (allowing billions of potential AVPs). This makes Diameter more flexible, allowing support for mobile remote users, for example.

TACACS AND TACACS+

The Terminal Access Controller Access Control System (**TACACS**) is a centralized access control system that requires users to send an ID and a static (reusable) password for authentication. TACACS uses UDP port 49 (and may also use TCP). Reusable passwords are a security vulnerability: The improved TACACS+ provides better password protection by allowing two-factor strong authentication.

TACACS+ is not backward compatible with TACACS. It uses TCP port 49 for authentication with the TACACS+ server.

PAP AND CHAP

The **Password Authentication Protocol (PAP)** is defined by RFC 1334 (<http://tools.ietf.org/html/rfc1334#section-2>) and is referred to as “not a strong authentication method.”¹ A user enters a password, which is sent across the network in clear text. Sniffing the network may disclose plaintext passwords.

The **Challenge Handshake Authentication Protocol (CHAP)**, defined by RFC 1994 (www.faqs.org/rfcs/rfc1994.html), provides protection against playback attacks.² It uses a central location that challenges remote users. As stated in the RFC,

CHAP depends upon a “secret” known only to the authenticator and the peer. The secret is not sent over the link. Although the authentication is only one-way, by negotiating CHAP in both directions the same secret set may easily be used for mutual authentication.³

PROCEDURAL ISSUES FOR ACCESS CONTROL

The day-to-day management of access control requires management of labels, clearances, formal access approval, and need to know. These formal mechanisms are typically used to protect highly sensitive data, such as for the government or the military.

Labels

Objects have labels, and as we will see in the next section, subjects have clearances. The object labels used by many world governments are confidential, secret, and top secret.

Additional labels exist, such as unclassified (data that is not sensitive), SBU (sensitive but unclassified) and for official use only (FOUO). SBU describes sensitive data that is not a matter of national security, such as the healthcare records of military enlisted personnel. This data must be protected, even though its release would not normally cause national security issues.

Private-sector companies use labels such as “Internal Use Only” and “Company Proprietary.”

Clearance

A **clearance** is a determination, typically made by a senior security professional, of whether or not a user can be trusted with a specific level of information. Clearances must determine the subject’s current and potential future trustworthiness—the latter is harder (and more expensive) to assess.

Rule-based access controls

Rule-based access control (RBAC) uses a series of defined rules, restrictions, and filters for accessing objects within a system. The rules are in the form of “if/then” statements. An example of a rule-based access control device is a proxy firewall that only allows users to web-surf to predefined approved content (e.g., “If the user is authorized to surf the Web, and the site is on the approved list, then allow access”).

Access control lists

Access control lists (ACLs) are used in many IT security policies, procedures, and technologies. They contain objects, and each entry describes the subjects that may access that object. Any attempt by a subject to access an object that does not have a matching entry in the ACL will be denied.

ACCESS CONTROL DEFENSIVE CATEGORIES AND TYPES

In order to understand and appropriately implement access controls, understanding the benefits that each control can add to security is vital.

There are six access control types:

- Preventive
- Detective
- Corrective
- Recovery
- Deterrent
- Compensating

Fast Facts

Access control types can fall into one of three categories: administrative, technical, or physical.

- **Administrative** (also called *directive*) controls are implemented by creating and following organizational policy, procedure, or regulation. User training and awareness fall into this category.
- **Technical** controls are implemented using software, hardware, or firmware that restricts logical access in an information technology system. Examples are firewalls, routers, encryption, and the like.
- **Physical** controls are implemented with physical devices, such as locks, fences, gates, and security guards.

Preventive

Preventive controls prevent actions. They apply restrictions to what a potential user, either authorized or unauthorized, can do. An example of an administrative preventive control is pre-employment drug screening, which is designed to prevent an organization from hiring an employee who is using illegal drugs.

Detective

Detective controls are controls that send alerts during or after an attack. Intrusion detection systems alerting after an attack, Closed-Circuit Television (CCTV) cameras alerting guards to an intruder, and a building alarm system that is triggered by an intruder are all examples of detective controls.

Corrective

Corrective controls “correct” a damaged system or process. They typically work hand in hand with detective controls. Antivirus software has both components. First, it runs a scan and uses its definition file to detect if there is any software that matches its virus list. If it detects a virus, the corrective controls take over, placing the suspicious software in quarantine or deleting it from the system.

Recovery

After a security incident has occurred, **recovery controls** may be needed to restore functionality to the system and the organization. Recovery means that the system must be recovered: reinstalled from OS media or images, data restored from backups, and so forth.

Deterrent

Deterrent controls deter users from performing actions on a system. Examples include a “beware of dog” sign: A thief facing two buildings, one with guard dogs and one without, is more likely to attack the building without.

Compensating

A **compensating** control is an additional security control put in place to compensate for weaknesses in others.

AUTHENTICATION METHODS

A key concept for implementing any type of access control is controlling the proper authentication of subjects within the IT system.

DID YOU KNOW?

There are three basic authentication methods: **Type 1** (something you know), **Type 2** (something you have), and **Type 3** (something you are). A fourth type of authentication is some place you are.

Type 1 Authentication: Something you know

Type 1 Authentication (something you know) requires testing the subject with some sort of challenge and response where the subject must respond with a knowledgeable answer. The subject is granted access based on something he or she knows, like a password or PIN (**Personal Identification Number**). This is the easiest, and often weakest, form of authentication.

PASSWORDS

Passwords have been the cornerstone of access control in IT systems. They are relatively easy and cheap to implement. There are four types of passwords to consider when implementing access controls: static, passphrases, one-time, and dynamic.

- **Static passwords** are reusable and may or may not expire. They are typically user generated and work best when combined with another authentication type, like a smart card or biometric control.
- **Passphrases** are long static passwords comprising words in a phrase or sentence. An example of a passphrase is “I will pass the CISSP in 6 months!” Passphrases may be made stronger by using nonsense words (replacing “CISSP” with “XYZZY,” for example), mixing case, and using additional numbers and symbols.
- **One-time passwords** may be used for a single authentication. They are very secure but difficult to manage. A one-time password is impossible to reuse and is valid for a one-time use only.
- **Dynamic passwords** change at regular intervals. RSA Security makes a synchronous token device called SecurID that generates a new token code every 60 seconds. The user combines his or her static PIN with the RSA dynamic token code to create one dynamic password that changes every time it is used. One drawback of dynamic passwords is the expense of the tokens themselves.

Strong authentication (also called *multi-factor authentication*) requires that the user present more than one authentication factor. For example, a user must insert an ATM card in order to withdraw money out of the bank but must also input the correct PIN.

Type 2 Authentication: Something you have

Type 2 authentication (something you have) requires that users possess something, such as a token, that proves they are authenticated. A token is a physical object that helps prove an identity claim.

SYNCHRONOUS DYNAMIC TOKEN

Synchronous Dynamic Tokens use time or counters to synchronize a displayed token code with the code expected by the authentication server: The codes are synchronized.

Time-based synchronous dynamic tokens display dynamic token codes that change frequently, such as every 60 seconds. The dynamic code is only good during that window. The authentication server knows the serial number of each authorized token, the user it is associated with, and the time. It can predict the dynamic code on each token using these three pieces of information.

Counter-based synchronous dynamic tokens use a simple counter: The authentication server expects token code 1, and the user’s token displays it. Once that code is used, the token displays the second code, so the server also expects token 2.

ASYNCHRONOUS DYNAMIC TOKEN

Asynchronous Dynamic Tokens are not synchronized with a central server. Their most common variety is challenge-response. Challenge-response token authentication systems produce a challenge, or input for the token device. The user then manually enters the information into the device along with her PIN, and the device produces an output. This output is sent to the system.

Type 3 Authentication: Something you are

Type 3 authentication (something you are) is biometrics, which uses physical characteristics as a means of identification or authentication. Biometrics can establish an identity or authenticate (prove an identity claim). For example: An airport facial recognition system can establish the identity of a known terrorist, and a fingerprint scanner can authenticate the identity of a subject (who makes the identity claim and then swipes his finger to prove it).

BIOMETRIC ENROLLMENT AND THROUGHPUT

Enrollment is the process of creating an account for the first time with a biometric system. Users typically provide their username (identity), a password or PIN, and provide biometric information such as by swiping their fingerprints on a fingerprint reader or having a photograph taken of their irises. Enrollment is a one-time process that should take two minutes or less.

Throughput is the process of authenticating to a biometric system. This is also called the biometric system response time. A typical throughput is six to ten seconds.

ACCURACY OF BIOMETRIC SYSTEMS

The accuracy of biometric systems should be considered before implementing a biometric control program. Three metrics are used to judge biometric accuracy: **False Reject Rate (FRR)**, **False Accept Rate (FAR)**, and **Crossover Error Rate (CER)**.

False Reject Rate

A false rejection occurs when an authorized subject is rejected by the biometric system as unauthorized. A false rejection is also called a **Type I error**. False rejections cause frustration in authorized users, reduction in work due to poor access conditions, and expenditure of resources to revalidate authorized users.

False Accept Rate

A false acceptance occurs when an unauthorized subject is accepted as valid. If an organization's biometric control is producing many false rejections, the overall control might have to lower the accuracy of the system by lessening the amount of data it collects when authenticating subjects. When the data points are lowered, the organization risks an increase in false acceptance rates, thus risking an unauthorized user gaining access. This type of error is also called a **Type II error**.

Crossover Error Rate

The crossover error rate describes the point where the False Reject Rate and False Accept Rate are equal. The CER is also known as the Equal Error Rate (EER). It describes the overall accuracy of a biometric system.

As the accuracy of a biometric system increases, FARs rise and FRRs drop. Conversely, as the accuracy decreases, FARs drop and FRRs rise. Figure 2.2 shows a graph depicting the FAR versus the FRR. The CER is the intersection of both lines of the graph, based on the 2007 IS Auditing Guideline: G36 Biometric Controls.⁴

TYPES OF BIOMETRIC CONTROL

A number of biometric controls are used today. The following subsections describe the major implementations and their specific pros and cons with regard to access control security.

Fingerprints

Fingerprints are the most widely used biometric control now available. The data used for storing each person's fingerprint must be small enough to be used for authentication. This data is a mathematical representation of fingerprint **minutiae**—that is, specific details of fingerprint friction ridges such as whorls and bifurcation. Figure 2.3 shows minutiae types (from left): bifurcation, ridge ending, core, and delta.⁵

Retina scan

A **retina scan** is a laser scan of the capillaries that feed the retina at the back of the eye. This can seem personally intrusive because the light beam must directly enter the pupil, and the user usually needs to press her eye up to an eye cup so that the scan can map the blood vessels of the retina. Information about the user's health can be gained through a retina scan: Conditions such as pregnancy and diabetes can be determined, which may raise legitimate privacy issues.

Iris scan

An **iris scan** is a passive biometric control. A camera takes a picture of the iris (the colored portion of the eye) and then compares it to photos within the authentication database. This also works via contact lenses and glasses.

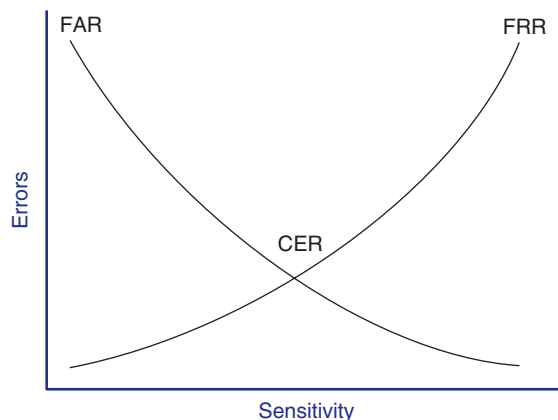


FIGURE 2.2

Crossover error rate.

Source: ISACA IS Auditing Guideline: G36 Biometric Controls. URL: www.isaca.org/Knowledge-Center/Standards/Pages/IS-Auditing-Guideline-G36-Biometric-Controls.aspx (accessed July 21, 2010).



FIGURE 2.3

Fingerprint minutiae.

Source: NIST Tech Beat, March 16, 2006. URL: www.nist.gov/public_affairs/techbeat/tb2006_0316.htm (accessed July 21, 2010).

Exam Warning

Retina scans are rarely used because of health risks and privacy issues. Alternatives should be considered for any biometric controls that risk exchange of bodily fluid or that raise legitimate privacy concerns.

Each person's two irises are unique, as are twins' irises. The benefits of iris scans include high accuracy, passive scanning (which may be accomplished without the subject's knowledge), and no exchange of bodily fluids.

Hand geometry

In **hand geometry** biometric control, measurements are taken from specific points on the subject's hand: "The devices use a simple concept of measuring and recording the length, width, thickness, and surface area of an individual's hand while guided on a plate."⁶ Hand geometry devices are fairly simple and can store information in as little as nine bytes.

Keyboard dynamics

Keyboard dynamics measure how hard a person presses each key and the rhythm with which the keys are pressed. As people learn how to type and use a computer keyboard, they develop specific habits that are difficult to impersonate, although not impossible.

Dynamic signature

Dynamic signatures measure the process by which someone signs his name. They are similar to keyboard dynamics, except that the method involved measures the handwriting of the subject while he signs his name.

Voice print

A **voice print** measures the subject's tone of voice while speaking a specific sentence or phrase. This type of access control is vulnerable to replay attacks (replaying a recorded voice), so other access controls must be implemented along with it. One such control requires subjects to say random words, protecting against an attacker playing prerecorded specific phrases. Another issue is that voices may substantially change because of illness, resulting in a false rejection.

Facial scan

Facial scan technology has greatly improved over the last few years. Also called facial recognition, it is the process of passively taking a picture of a subject's face and comparing that picture to a list stored in a database.

Someplace you are

Someplace you are describes location-based access control using technologies such as the global positioning system (GPS), IP address-based geo-location, or the physical location for a point-of-sale purchase. These controls can deny access if the subject is not in the right place.

ACCESS CONTROL TECHNOLOGIES

Several technologies are used for the implementation of access controls. As each technology is presented, it is important to identify what is unique about it.

Single Sign-On

In Single Sign-On (SSO), multiple systems use a central authentication server (AS). This allows users to authenticate once and then access multiple different systems. It also allows security administrators to add, change, or revoke user privileges on one central system.

The primary disadvantage of SSO is that it may allow an attacker to gain access to multiple resources after compromising one authentication method such as a password.

Kerberos

Kerberos is a third-party authentication service that supports Single Sign-On (www.kerberos.org/) Kerberos (or Cerberus) was the name of the three-headed dog that guarded the entrance to Hades in Greek mythology.

KERBEROS CHARACTERISTICS

Kerberos uses secret key encryption and provides mutual authentication of both clients and servers. It protects against network sniffing and replay attacks. The current version of Kerberos is version 5, described by RFC 4120 (www.ietf.org/rfc/rfc4120.txt).

Fast Facts

Kerberos has the following components:

- **Principal:** client (user) or service
- **Realm:** logical Kerberos network
- **Ticket:** data that authenticates a principal's identity
- **Credentials:** a ticket and a service key
- **KDC:** Key Distribution Center, which authenticates principals
- **TGS:** Ticket Granting Service
- **TGT:** Ticket Granting Ticket
- **C/S:** Client/Server, regarding communications between the two

KERBEROS OPERATIONAL STEPS

A Kerberos principal, a client run by user Alice, wishes to access a printer. Alice may print after taking these five (simplified) steps:

1. Kerberos Principal Alice contacts the KDC (Key Distribution Center, which acts as an authentication server), requesting authentication.
2. The KDC sends Alice a session key, encrypted with Alice's secret key. The KDC also sends a TGT (Ticket Granting Ticket), encrypted with the TGS's secret key
3. Alice decrypts the session key and uses it to request permission to print from the TGS
4. Seeing that Alice has a valid session key (and therefore has proven her identity claim), the TGS sends her a C/S session key (second session key) to use to print. The TGS also sends a service ticket encrypted with the printer's key.
5. Alice connects to the printer. The printer, seeing a valid C/S session key, knows she has permission to print, and also knows that she is authentic

This process is summarized in Figure 2.4.

The session key in step 2 of Figure 2.4 is encrypted with Alice's key (represented as $\{\text{Session Key}\}_{\text{Key}^{\text{Alice}}}$). Note that the TGT is encrypted with the TGS's key: Alice cannot decrypt the TGT (only the TGS can); she simply sends it to the TGS. The TGT contains a number of items, including a copy of Alice's session key. This is how the TGS knows that Alice has a valid session key (which proves that Alice is authenticated).

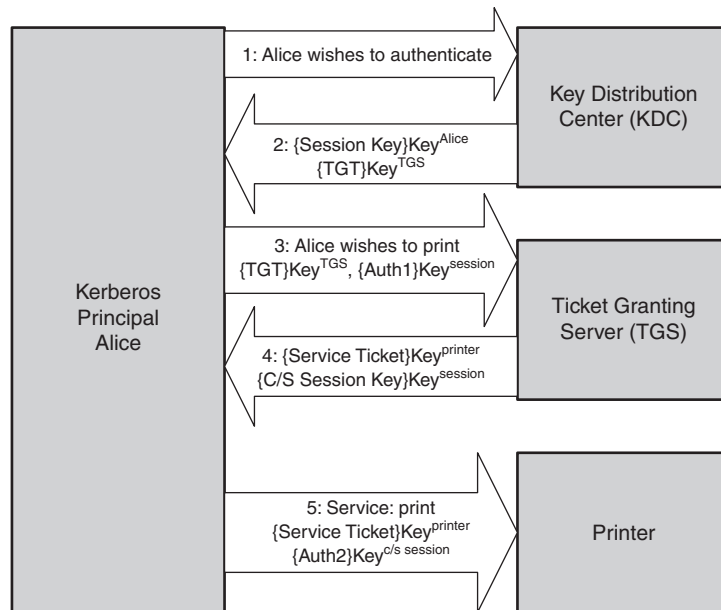


FIGURE 2.4

Kerberos steps.

Source: Kerberos. URL: www.giac.org/resources/whitepaper/cryptography/47.php (accessed July 21, 2010).

The TGT is good for a site-selected specific lifetime, often set at 10 hours. This allows a typical user to authenticate once and access network resources for the lifetime of the ticket. Kerberos is stateless for this reason: Once Alice has a TGT, she may use it for its lifetime, even if the KDC goes offline. Also, the TGS can allow Alice to print without consulting the KDC: Everything the TGS needs to know is contained in the traffic Alice sends, including the TGT and the first authenticator.

The same is true for the service ticket Alice sends to the printer. It is encrypted with the printer's key and contains a copy of the client/server session key. Alice cannot decrypt it and simply passes it back to the printer. This allows the printer to make its decision based entirely on what Alice sends, without consulting the KDC or the TGS.

SESAME

The Secure European System for Applications in a Multi-vendor Environment (SESAME) is a single sign-on system that supports heterogeneous environments. It can be thought of as a sequel of sorts to Kerberos: "SESAME adds to Kerberos: heterogeneity, sophisticated access control features, scalability of public key systems, better manageability, audit and delegation."⁷ Of those improvements, the addition of public key (asymmetric) encryption is the most compelling because it addresses one of the biggest weaknesses in Kerberos: the plaintext storage of symmetric keys.

SESAME uses Privilege Attribute Certificates (PACs) in place of Kerberos tickets. More information on SESAME is available at www.cosic.esat.kuleuven.be/sesame/.

ASSESSING ACCESS CONTROL

A number of processes exist to assess the effectiveness of access control. Tests with a narrow scope include penetration, vulnerability assessments, and security audits. A security assessment is a broader test that may include the narrower tests (e.g., penetration), as subsections.

Penetration testing

A penetration tester is a white-hat hacker who receives authorization to attempt to break into an organization's physical or electronic perimeter (and sometimes both). **Penetration tests** ("pen tests" for short) are designed to determine whether black-hat hackers can do the same. They are narrow, but often useful, especially if the penetration tester is successful.

Penetration tests may include the following:

- Network (Internet)
- Network (internal or DMZ)
- War dialing

- Wireless
- Physical (attempt to gain entrance to a facility or room)
- Wireless

Network attacks may leverage client-side, server-side, or web application attacks. (See Chapter 5 for more information on these.) **War dialing** uses a modem to dial a series of phone numbers, looking for an answering modem carrier tone (the penetration tester then attempts to access the answering system); the name derives from the 1983 movie *WarGames*.

Social engineering uses the human mind to bypass security controls. It may be used in combination with many types of attacks, especially client-side attacks or physical tests

A **zero knowledge** (or *black-box*) test is “blind”; the penetration tester starts with no external or trusted information and begins with public information only. A **full knowledge test** (also called *crystal-box*) provides internal information to the penetration tester, including network diagrams, policies and procedures, and sometime reports from previous testers. **Partial knowledge** tests are between zero and full knowledge: The penetration tester receives limited trusted information.

Vulnerability testing

Vulnerability scanning (also called vulnerability testing) scans a network or system for a list of predefined vulnerabilities such as system misconfiguration, outdated software, or a lack of patching. A vulnerability testing tool such as Nessus (www.nessus.org) or OpenVAS (www.openvas.org) may be used to identify the vulnerabilities.

Security audits

A **security audit** is a test against a published standard. Organizations may be audited for PCI (Payment Card Industry) compliance, for example. PCI includes many required controls such as firewalls, specific access control models, and wireless encryption. An auditor verifies that a site or organization meets the published standard.

Security assessments

Security assessments are a holistic approach to assessing the effectiveness of access control. Instead of looking narrowly at penetration tests or vulnerability assessments, security assessments have a broader scope.

SUMMARY OF EXAM OBJECTIVES

If one thinks of the castle analogy for security, access control would be the moat and the castle walls. It ensures that the border protection mechanisms, from both a logical and a physical viewpoint, are secured. Its purpose is to

allow authorized users access to appropriate data and to deny access to unauthorized users. This is also known as limiting subjects' access to objects. Even though this task is complex and involved, it is possible to implement a strong access control program without overburdening the users who rely on access to the system.

Protecting the CIA triad is another key aspect to implementing access controls. Maintaining confidentiality, integrity, and availability is of the utmost importance. Maintaining the security of a CIA's system means enacting specific procedures for data access. These procedures will change depending on the functionality that users require and the sensitivity of the data stored in the system.

TOP FIVE TOUGHEST QUESTIONS

1. Which group launches the most attacks?
 - A. Insiders
 - B. Outsiders
 - C. Hacktivists
 - D. Script kiddies

Questions 2 and 3 are based on this scenario: Your company has hired a third party to conduct a penetration test. Your CIO wants to know if exploitation of critical business systems is possible. The following are the company's two requirements:

- The tests will be conducted on live business-functional networks. These networks must be functional in order for business to run and cannot be shut down, even for an evaluation.
 - The test will be as in depth as possible.
2. Which kind of test should be recommended?
 - A. Zero knowledge
 - B. Partial knowledge
 - C. Full knowledge
 - D. Vulnerability
 3. While conducting the penetration test, the tester discovers that a critical business system is currently compromised. What should he do?
 - A. Note the results in the penetration testing report
 - B. Immediately end the penetration test and call the CIO
 - C. Remove the malware
 - D. Shut the system down
 4. A policy that a user must have a business requirement to view data before attempting to do so is an example of enforcing what?
 - A. Least privilege
 - B. Need to know
 - C. Rotation of duties
 - D. Separation of duties

5. Which technique would raise the False Accept Rate (FAR) and lower the False Reject Rate (FRR) in a fingerprint-scanning system?
- A. Decrease the amount of minutiae that is verified
 - B. Increase the amount of minutiae that is verified
 - C. Lengthen the enrollment time
 - D. Lower the throughput time

Answers

1. Correct Answer and Explanation: **B**. Answer **B** is correct; outsiders launch most attacks (though most fail).
Incorrect Answers and Explanations: **A**, **C**, and **D**. Answers **A**, **C**, and **D** are incorrect. Insiders may launch the most successful attacks that cause the highest impact, but most attacks are launched from the outside (and typically mitigated). Hacktivists and Script kiddies are usually subsets of outsiders, making **B** the best answer.
2. Correct Answer and Explanation: **C**. **C** is the correct answer; the customer wants a full evaluation but is worried because of the importance of the network. Because the customer wants as full an evaluation as possible but does not want the network in any kind of jeopardy, a full knowledge assessment is necessary because it is the only one that will allow for the most in-depth analysis with the least amount of risk to the network.
Incorrect Answers and Explanations: **A**, **B**, and **D**. **A** is incorrect because a zero knowledge test will not produce the most in-depth assessment of the network. **B** is incorrect because partial knowledge, although better than zero knowledge, will not produce the necessary assessment. **D** is incorrect because vulnerability testing does not exploit systems, which is a requirement of the test.
3. Correct Answer and Explanation: **B**. Answer **B** is correct; when discovering a live malicious intrusion, the penetration tester should immediately end the test and notify the client.
Incorrect Answers and Explanations: **A**, **C**, and **D**. Answers **A**, **C**, and **D** are incorrect. Noting the results is not enough: System integrity, data integrity, and confidentiality are compromised or at risk and immediate action is required. Removing the malware may cause more damage and/or alert the attackers to the penetration tester's presence. Attackers may become more malicious if they believe they have been discovered. Shutting the system down will harm availability (and possibly integrity), and will destroy any evidence that exists in memory.
4. Correct Answer and Explanation: **B**. Answer **B** is correct; need to know means that the user must have a need (requirement) to access a specific object before doing so.
Incorrect Answers and Explanations: **A**, **C**, and **D**. Answers **A**, **C**, and **D** are incorrect. Least privilege is less granular than need to know: Users have the least amount of privilege to do their jobs, but objects are still typically grouped together (such as for allowing access to all backup tapes for a

backup administrator). Separation of duties is designed to divide sensitive tasks among multiple subjects. Rotation of duties is designed to mitigate collusion.

5. Correct Answer and Explanation: **A**. Answer **A** is correct; decreasing the amount of minutiae will lower the accuracy of the system, which will lower false rejects but raise false accepts.

Incorrect Answers and Explanations: **B**, **C**, and **D**. Answers **B**, **C**, and **D** are incorrect. Increasing the amount of minutiae will make the system more accurate, increasing the FRR and lowering the FAR. Enrollment and throughput time are not directly connected to FAR and FRR.

Endnotes

1. RFC 1334 PAP. URL: <http://tools.ietf.org/html/rfc1334#section-2> (accessed July 21, 2010).
2. RFC 1994 CHAP. URL: www.faqs.org/rfcs/rfc1994.htm1 (accessed July 21, 2010).
3. Ibid.
4. ISACA IS Auditing Guideline: G36 Biometric Controls. URL: www.isaca.org/Knowledge-Center/Standards/Pages/IS-Auditing-Guideline-G36-Biometric-Controls.aspx (accessed July 21, 2010).
5. NIST Tech Beat, March 16, 2006. URL: www.nist.gov/public_affairs/techbeat/tb2006_0316.htm (accessed July 21, 2010).
6. Hand Geometry. URL: www.biometrics.gov/Documents/HandGeometry.pdf (accessed July 21, 2010).
7. SESAME in a Nutshell. URL: www.cosic.esat.kuleuven.be/sesame/html/sesame_what.html (accessed July 21, 2010).

This page intentionally left blank

CHAPTER 3

Domain 3: Cryptography

39

Exam Objectives in this Chapter

- Cornerstone cryptographic concepts
- Symmetric encryption
- Asymmetric encryption
- Hash functions
- Cryptographic attacks
- Implementing cryptography

INTRODUCTION

Cryptography is secret writing: secure communication that can be understood only by the intended recipient. While the fact that data is being transmitted may be known, the content of that data should remain unknown to third parties. Data in motion (moving on a network) and at rest (stored on a device such as a disk) may be encrypted.

CORNERSTONE CRYPTOGRAPHIC CONCEPTS

Fundamental cryptographic concepts are embodied by all strong encryption. They must be understood before learning about specific implementations.

Key terms

Cryptology is the science of secure communications. **Cryptography** creates messages whose meaning is hidden. **Cryptanalysis** is the science of breaking encrypted messages (recovering their meaning). Many use the term *cryptography* in place of *cryptology*; however, it is important to remember that cryptology encompasses both cryptography and cryptanalysis.

A **cipher** is a cryptographic algorithm. A **plaintext** message is an unencrypted message. **Encryption** converts the plaintext to **ciphertext**. **Decryption** turns ciphertext back into plaintext.

Confidentiality, integrity, authentication, and nonrepudiation

Cryptography can provide confidentiality (secrets remain secret) and integrity (data is not altered in an unauthorized manner): it's important to note that it does not directly provide availability. Cryptography can also provide authentication (proving an identity claim).

Additionally, cryptography can provide **nonrepudiation**, which is assurance that a specific user performed a specific transaction *and* assurance that the transaction did not change.

Substitution and permutation

Cryptographic substitution replaces one character for another. Permutation (also called *transposition*) rearranges the characters of the plaintext anagram-style. "ATTACKATDAWN" can be rearranged to "CAAKDTANTATW," for example. Substitution and permutation are often combined.

DID YOU KNOW?

Strong encryption destroys patterns. If a single bit of plaintext changes, the odds of every bit of resulting ciphertext changing should be 50/50. Any signs of nonrandomness may be used as clues by a cryptanalyst, hinting at the underlying order of the original plaintext or key.

Cryptographic strength

Good encryption is strong: For key-based encryption, it should be very difficult (and ideally impossible) to convert ciphertext back to plaintext without the key. The work factor describes how long it will take to break a cryptosystem (i.e., to decrypt a ciphertext without the key).

Secrecy of the cryptographic algorithm does not provide strength. In fact, secret algorithms are often proven to be quite weak. Strong crypto relies on math, not secrecy, to provide strength. Ciphers that have stood the test of time are public algorithms such as the Triple Data Encryption Standard (TDES) and the Advanced Encryption Standard (AES).

Monoalphabetic and polyalphabetic ciphers

A **monoalphabetic cipher** uses one alphabet: A specific letter (say *E*) is substituted for another (say *X*). A **polyalphabetic cipher** uses multiple alphabets: *E* may be substituted for *X* in one round and then *S* in the next round. Monoalphabetic ciphers are susceptible to frequency analysis.

Exclusive Or

Exclusive Or (XOR) is the "secret sauce" behind modern encryption. Combining a key with a plaintext via XOR creates ciphertext. XORing the same key to the ciphertext restores the original plaintext.

Table 3.1 XOR Truth Table

X	Y	X XOR Y
0	0	0
0	1	1
1	0	1
1	1	0

Two bits are true (or 1) if one or the other (exclusively, not both) is 1. If both bits are 0 or both bits are 1, they XOR to 0. XOR uses a truth table (see Table 3.1). In the truth table, a 0 is false and a 1 is true. This dictates how to combine the bits of a key and plaintext.

Types of cryptography

There are three primary types of modern encryption: **symmetric**, **asymmetric**, and **hashing**, which will be discussed in the following sections. Symmetric encryption uses one key: The same key encrypts and decrypts. Asymmetric cryptography uses two keys: If you encrypt with one key, you can decrypt with the other. Hashing is a one-way cryptographic transformation using an algorithm (and no key).

SYMMETRIC ENCRYPTION

Symmetric encryption uses one key to encrypt and decrypt. If you encrypt a zip file and then decrypt it with the same key, you are using symmetric encryption. In symmetric encryption, also called “secret key” encryption, the key must be kept secret from third parties. Strengths include speed and cryptographic strength per bit of key. The major weakness is that the key must be securely shared before two parties may securely communicate. Symmetric keys are often shared via an out-of-band method, such as face-to-face discussion.

Stream and block ciphers

Symmetric encryption may have stream and block modes. Stream mode means that each bit is independently encrypted in a “stream.” Block mode ciphers encrypt blocks of data each round: for example, 56 bits for the Data Encryption Standard (DES) and 128, 192, or 256 bits for AES. Some block ciphers can emulate stream ciphers by setting the block size to 1 bit.

DES

DES is the Data Encryption Standard, which describes the Data Encryption Algorithm (DEA). DES was designed by IBM, based on their older Lucifer symmetric cipher. It uses a 64-bit block size (meaning that it encrypts 64 bits each round) and a 56-bit key.

Exam Warning

Even though “DES” is commonly referred to as an algorithm, it is technically the name of the published standard that describes DEA. It may sound like splitting hairs, but that is an important distinction to keep in mind on the exam. “DEA” may be the best answer for a question regarding the algorithm itself.

MODES OF DES

DES can use five different modes to encrypt data. The modes’ primary differences are block versus (emulated) stream, the use of initialization vectors, and whether errors in encryption will propagate to subsequent blocks.

Fast Facts

The five modes of DES are

- Electronic Code Book (ECB)
- Cipher Block Chaining (CBC)
- Cipher Feedback (CFB)
- Output Feedback (OFB)
- Counter (CTR)

ECB is the original mode of DES. CBC, CFB, and OFB were later added in Federal Information Processing Standards (FIPS) Publication 81 (see www.itl.nist.gov/fipspubs/fip81.htm). CTR is the newest mode, described in NIST Special Publication 800-38a (see <http://csrc.nist.gov/publications/nistpubs/800-38a/sp800-38a.pdf>).

Electronic Code Book

Electronic Code Book (ECB) is the simplest and weakest form of DES. It uses no initialization vector or chaining. Identical plaintexts with identical keys encrypt to identical ciphertexts. Two plaintexts with partial identical portions (such as the header of a letter) encrypted with the same key have partial identical ciphertext portions. ECB may also leave plaintext patterns evident in the resulting ciphertext.

Cipher Block Chaining

Cipher Block Chaining (CBC) is a block mode of DES that XORs the previous encrypted block of ciphertext to the next block of plaintext to be encrypted. The first encrypted block is an initialization vector that contains random data. This “chaining” destroys patterns. One limitation of CBC is that encryption errors propagate: An encryption error in one block will cascade through subsequent blocks because of the chaining, thus destroying their integrity.

Cipher Feedback

Cipher Feedback (CFB) is very similar to CBC; the primary difference is that CFB is a stream mode, using feedback to destroy patterns. Like CBC, CFB uses an initialization vector; patterns are destroyed and errors propagate.

Output Feedback

Output Feedback (OFB) differs from CFB in the way feedback is accomplished: CFB uses the previous ciphertext—that is, the previous ciphertext is the subkey XORed to the plaintext. OFB uses the subkey *before* it is XORed to the plaintext. Since the subkey is not affected by encryption errors, errors do not propagate.

Counter

Counter (CTR) is like OFB; the difference again is the feedback: CTR uses a counter. This mode shares the same advantages as OFB (patterns are destroyed and errors do not propagate) with an additional advantage: since the feedback can be as simple as an ascending number, CTR encryption can be done in parallel. A simple example would be the first block is XORed to the number 1, the second to the number 2, and so forth. Any number of rounds can be combined in parallel in this way.

Table 3.2 summarizes the 5 modes of DES.

SINGLE DES

Single DES is the original DES implementation, encrypting 64-bit blocks of data with a 56-bit key using 16 rounds of encryption. The work factor required to break DES was reasonable in 1976, but advances in CPU speed and parallel architecture have made it vulnerable to a brute-force key attack today, where every possible key is generated and attempted.

TRIPLE DES

Triple DES applies single DES encryption three times per block. Formally called the Triple Data Encryption Algorithm (TDEA) and commonly called TDES, it became a recommended standard in 1999.

Table 3.2 Modes of DES

Mode	Type	Initialization vector	Error propagation?
Electronic Code Book (ECB)	Block	No	No
Cipher Block Chaining (CBC)	Block	Yes	Yes
Cipher Feedback (CFB)	Stream	Yes	Yes
Output Feedback (OFB)	Stream	Yes	No
Counter Mode (CTR)	Stream	Yes	No

Triple DES encryption order and keying options

Triple DES applies DES encryption three times per block. FIPS 46-3 describes this as “Encrypt, Decrypt, Encrypt” (EDE) order using three keying options: one, two, or three unique keys (1TDES EDE, 2TDES EDE, and 3TDES EDE, respectively).

International Data Encryption Algorithm

The International Data Encryption Algorithm is a symmetric block cipher designed as an international replacement for DES. Patented in many countries, IDEA uses a 128-bit key and a 64-bit block size.

Advanced Encryption Standard

The Advanced Encryption Standard (AES) is the current U.S. standard symmetric block cipher. It uses 128-bit keys (10 rounds of encryption), 192-bit keys (12 rounds of encryption), or 256-bit keys (14 rounds of encryption) to encrypt 128-bit blocks of data.

CHOOSING AES

The U.S. National Institute of Standards and Technology (NIST) solicited input on a replacement for DES in the *Federal Register* in January 1997. Fifteen AES candidates were announced in August 1998, and the list was reduced to five in August 1999. Table 3.3 lists the five AES finalists. Rijndael was chosen and became AES, which has four functions: SubBytes, ShiftRows, MixColumns, and AddRoundKey.

Blowfish and Twofish

Blowfish and Twofish are symmetric block ciphers created by teams lead by Bruce Schneier, author of *Applied Cryptography*. Blowfish uses 32-bit through 448-bit keys (the default is 128) to encrypt 64 bits of data. Twofish was an AES finalist, encrypting 128-bit blocks using 128-bit through 256-bit keys. Both are open algorithms, unpatented and freely available.

RC5 and RC6

RC5 and RC6 are symmetric block ciphers created by RSA Laboratories. RC5 uses 32-bit (testing purposes), 64-bit (replacement for DES), or 128-bit blocks. The key size ranges from 0 to 2040 bits.

Table 3.3 Five AES Finalists

Name	Author
MARS	IBM (11 authors)
RC6	RSA (Rivest, Robshaw, Sidney, Yin)
Rijndael	Daemen, Rijmen
Serpent	Anderson, Biham, Knudsen
Twofish	Schneier, Kelsey, Hall, Ferguson, Whiting, Wagner

RC6 was an AES finalist. It is based on RC5 but altered to meet the AES requirements. It is also stronger than RC5, encrypting 128-bit blocks using 128-, 192-, or 256-bit keys.

ASYMMETRIC ENCRYPTION

Asymmetric encryption uses two keys: If you encrypt with one, you may decrypt with the other. One key may be made public (the **public key**), which is why asymmetric encryption is also called public key encryption. Anyone who wants to communicate with you may simply download your publicly posted key and use it to encrypt her plaintext. Once the plaintext is encrypted, your public key cannot decrypt it: Only your **private key** can do so. As the name implies, your private key must be kept private and secure.

Asymmetric methods

Math lies behind the asymmetric breakthrough. Mathematical methods use “one-way functions,” which are easy to compute “one way” and difficult to compute in the reverse direction.

FACTORING PRIME NUMBERS

An example of a one-way function is factoring a composite number into its primes. Multiplying the prime number 6,269 by the prime number 7,883 results in the composite number 49,418,527. That “way” is quite easy to compute, taking milliseconds on a calculator. Answering the question “Which prime number times which prime number equals 49,418,527” is *much* more difficult. This problem is called factoring, and no shortcut for it has been found in hundreds of years. It is the basis of the RSA algorithm.

DISCRETE LOGARITHM

A logarithm is the opposite of exponentiation. Computing 7 to the 13th power (exponentiation) is easy on a modern calculator: 96,889,010,407. Asking, “96,889,010,407 is what to what power?” (finding the logarithm) is more difficult. Discrete logarithms are applied to groups, which is a much harder problem. This one-way function is the basis of the Diffie-Hellman and ElGamal asymmetric algorithms.

Diffie-Hellman Key Agreement Protocol

Key agreement allows two parties to securely agree on a symmetric key via a public channel, such as the Internet, with no prior key exchange. An attacker who is able to sniff the entire conversation cannot derive the exchanged key. The Diffie-Hellman Key Agreement Protocol (also called the Diffie-Hellman Key Exchange) was created in 1976 by Whitfield Diffie and Martin Hellman. It uses discrete logarithms to provide security.

ELLIPTIC CURVE CRYPTOGRAPHY

Elliptic Curve Cryptography (ECC) leverages a one-way function that uses discrete logarithms as applied to elliptic curves. Solving this problem is harder than solving discrete logarithms, so algorithms based on ECC are much stronger per

bit than systems using discrete logarithms (and also stronger than factoring prime numbers). ECC requires fewer computational resources because shorter keys can be used compared with other asymmetric methods. It is often used in lower-power devices for this reason.

ASYMMETRIC AND SYMMETRIC TRADE-OFFS

Asymmetric encryption is far slower than symmetric encryption and is also weaker per bit of key length. The strength of asymmetric encryption is in its ability to securely communicate without presharing a key.

HASH FUNCTIONS

A hash function provides encryption using an algorithm and no key. It is referred to as “one way” because there is no way to reverse the encryption. A variable-length plaintext is “hashed” into a fixed-length hash value (often called a “message digest” or simply a “hash”).

Hash functions are primarily used to provide integrity: If the hash of a plaintext changes, the plaintext itself changes. Common older hash functions include Secure Hash Algorithm 1 (**SHA-1**), which creates a 160-bit hash, and Message Digest 5 (**MD5**), which creates a 128-bit hash. Weaknesses have been found in both MD5 and SHA-1; newer alternatives such as SHA-2 are recommended.

MD5

MD5, created by Ronald Rivest, is the most widely used of the MD family of hash algorithms. It creates a 128-bit hash value based on any input length. MD5 was quite popular over the years, but weaknesses have been discovered where collisions can be found in a practical amount of time. MD6 is the newest member of the MD family, first published in 2008.

Secure Hash Algorithm

Secure Hash Algorithm (SHA) is the name of a series of hash algorithms. SHA-1 creates a 160-bit hash value. SHA-2 includes SHA-224, SHA-256, SHA-384, and SHA-512, named after the length of the message digest each creates.

HAVAL

Hash of Variable Length (HAVAL) creates message digests of 128, 160, 192, 224, or 256 bits in length, in three, four, or five rounds. It uses some of the design principles behind the MD family and is faster than MD5.

CRYPTOGRAPHIC ATTACKS

Cryptographic attacks are used by cryptanalysts to recover plaintext without a key. Please remember that recovering the key (sometimes called “stealing the key”) is usually easier than breaking modern encryption.

Brute force

A brute-force attack generates the entire keyspace, which is every possible key. Given enough time, the plaintext will be recovered.

Known plaintext

A known plaintext attack relies on recovering and analyzing a matching plaintext and ciphertext pair: The goal is to derive the key that was used. You may be wondering why you would need the key if you already have the plaintext: Recovering the key allows you to decrypt other ciphertexts encrypted with the same key.

Chosen plaintext and adaptive chosen plaintext

A cryptanalyst chooses the plaintext to be encrypted in a plaintext attack; the goal is to derive the key. Encrypting without knowing the key is accomplished via an “encryption oracle,” or a device that encrypts without revealing the key.

Adaptive chosen plaintext begins with a chosen plaintext attack in round 1. The cryptanalyst then “adapts” further rounds of encryption based on the previous round.

Chosen ciphertext and adaptive chosen ciphertext

Chosen ciphertext attacks mirror chosen plaintext attacks. The difference is that the cryptanalyst chooses the ciphertext to be decrypted. This attack is usually launched against asymmetric cryptosystems, where the cryptanalyst may choose public documents to decrypt that are signed (encrypted) with a user’s public key.

Adaptive chosen ciphertext also mirrors its plaintext cousin: It begins with a chosen ciphertext attack in round 1. The cryptanalyst then “adapts” further rounds of decryption based on the previous round.

Meet-in-the-middle attack

A meet-in-the-middle attack encrypts on one side, decrypts on the other side, and meets in the middle. The most common attack is against “double DES,” which encrypts with two keys in “encrypt, encrypt” order. This is a known plaintext attack: The attacker has a copy of a matching plaintext and ciphertext, and seeks to recover the two keys used to encrypt.

Known key

The term “known key attack” is misleading: If the cryptanalyst knows the key, the attack is over. “Known key” means that the cryptanalyst knows something about the key, thus reducing the effort needed to attack it. If the cryptanalyst knows the key is uppercase letters and numbers only, other characters may be omitted in the attack.

Differential cryptanalysis

Differential cryptanalysis seeks to find the “difference” between related encrypted plaintexts, which may differ by a few bits. This attempt is usually launched as an adaptive chosen plaintext attack: The attacker chooses the plaintext to be encrypted (but does not know the key) and then encrypts related plaintexts.

Linear cryptanalysis

Linear cryptanalysis is a known plaintext attack where the cryptanalyst finds large amounts of plaintext/ciphertext pairs created with the same key. The pairs are studied to derive information about the key used to create them.

Both differential and linear analysis can be combined as differential linear analysis.

Side-channel attacks

Side-channel attacks use physical data to break a cryptosystem, such as by monitoring CPU cycles or power consumption used while encrypting or decrypting.

Birthday attack

The birthday attack is named after the birthday paradox. The name is based on the fact that in a room with 23 people or more, the odds are greater than 50% that two will share the same birthday. Many find this counterintuitive, and the birthday paradox illustrates why many people’s instinct in probability (and risk) is wrong. You are not trying to match a specific birthday (e.g., yours); you are trying to match any birthday.


IMPLEMENTING CRYPTOGRAPHY

The three types of cryptography—symmetric, asymmetric, and hash-based—do not exist in a vacuum; they are applied in the real world, often in combination, to provide confidentiality, integrity, authentication, and nonrepudiation.

Digital signatures


Digital signatures are used to cryptographically sign documents. They provide nonrepudiation, which includes authentication of the signer’s identity and proof of the document’s integrity (i.e., that the document did not change). This means that the sender cannot later deny (or repudiate) signing the document.

Let’s say that Roy wants to send a digitally signed email to Rick. He writes the email, which is the plaintext. He then uses the SHA-1 hash function to generate a hash value of the plaintext. Next he creates the digital signature by encrypting the hash with his RSA private key. (See Figure 3.1.) Finally, Roy attaches the signature to his plaintext email and hits send.

**FIGURE 3.1**

Creating a digital signature.

Source: Adapted from the film *Bladerunner*, directed by Ridley Scott, 1982.

**FIGURE 3.2**

Verifying a digital signature.

Source: Adapted from the film *Bladerunner*, directed by Ridley Scott, 1982.

Rick receives Roy's email and generates his own SHA-1 hash value of its plaintext. Next he decrypts the digital signature with Roy's RSA public key, recovering the SHA-1 hash Roy generated. Rick then compares his SHA-1 hash with Roy's. (See Figure 3.2.)

If the two hashes match, Rick knows two things:

1. Roy must have sent the email (only Roy knows his private key). This authenticates Roy as the sender.
2. The email did not change. This proves its integrity.

If the hashes match, Roy cannot later deny having signed the email. This is non-repudiation. If the hashes do not match, Rick knows either that Roy did not send it or that the email's integrity was violated.

HMAC

A Hashed Message Authentication Code (**HMAC**) combines symmetric encryption with hashing. The approach is similar to a digital signature, except it uses symmetric, instead of asymmetric, encryption. HMACs are used by IPsec (see the section IPsec to come).

CBC-MAC

Cipher Block Chaining Message Authentication Code (CBC-MAC) uses the CBC mode of a symmetric block cipher such as DES to create a message authentication code (MAC). A CBC-MAC provides integrity. This method differs from HMAC because it uses one algorithm; HMACs use two: a hash such as SHA-1 followed by a symmetric block cipher such as DES or AES.

Public Key Infrastructure

Public Key Infrastructure (PKI) leverages all three forms of encryption to provide and manage digital certificates. A digital certificate is a public key signed with a digital signature. It may be server-based (used for SSL websites such as *www.ebay.com*, for example) or client-based (bound to a person). If the two are used together, they provide mutual authentication and encryption. The standard digital certificate format is X.509.

CERTIFICATE AUTHORITIES

Digital certificates are issued by Certificate Authorities (CAs), which authenticate the identity of persons or organizations before issuing a certificate to them. CAs may be private (run internally) or public (such as Verisign or Thawte).

CERTIFICATE REVOCATION LISTS

The CAs maintain Certificate Revocation Lists (CRLs), which, as the name implies, list certificates that have been revoked. A certificate may be revoked if the private key has been stolen, an employee is terminated, and so forth.

IPsec

IPsec (Internet Protocol Security) is a suite of protocols that add a cryptographic layer to both IPv4 and IPv6. It is one of the methods for providing Virtual Private Networks (VPN), which allow you to send private data over an insecure network such as the Internet (the data crosses a public network, but is “virtually private”). IPsec includes two primary protocols: Authentication Header (AH) and Encapsulating Security Payload (ESP). AH and ESP provide different, and sometimes overlapping, functionality.

Supporting IPsec protocols include the Internet Security Association and Key Management Protocol (ISAKMP) and the Internet Key Exchange (IKE).

AH AND ESP

Authentication Header (AH) provides authentication and integrity for each packet of network data. It offers no confidentiality, but rather acts as a digital signature for the data. AH also protects against replay attacks, where data is sniffed off a network and resent, often in an attempt to fraudulently reuse encrypted authentication credentials.

Encapsulating Security Payload (ESP) primarily provides confidentiality by encrypting packet data. It may also optionally provide authentication and integrity.

SECURITY ASSOCIATION AND ISAKMP

AH and ESP may be used separately or in combination. An IPsec Security Association (SA) is a simplex (one-way) connection that may be used to negotiate ESP or AH parameters. If two systems communicate via ESP, they use two SAs (one for each direction). If the systems leverage AH in addition to ESP, they use two more SAs, for a total of four. Each simplex SA connection is identified by a unique 32-bit number known as the Security Parameter Index (SPI). The SA process is managed by the Internet Security Association and Key Management Protocol (ISAKMP).

TUNNEL AND TRANSPORT MODE

IPsec can be used in either tunnel mode or transport mode. Tunnel mode is used by security gateways (which can provide point-to-point IPsec tunnels). It encrypts the entire packet, including the original packet headers. ESP transport mode only encrypts the data (and not the original headers); this mode is commonly used when the sending and receiving system can “speak” IPsec natively.

Crunch Time

AH authenticates the original IP headers, so it is often used (along with ESP) in transport mode because the original headers are not encrypted. Tunnel mode typically

uses ESP alone (the original headers are encrypted, and thus protected, by ESP).

IKE

IPsec can employ a variety of encryption algorithms, such as MD5 or SHA-1 for integrity and triple DES or AES for confidentiality. The algorithm selection process is negotiated by the Internet Key Exchange (IKE). Two sides of an IPsec tunnel typically use IKE to negotiate to the highest and fastest level of security, for example selecting AES over single DES for confidentiality if both sides support it.

SSL and TLS

Secure Sockets Layer (SSL) brought the power of PKI to the Web. It authenticates and provides confidentiality to web traffic. Transport Layer Security (TLS) is the successor to SSL. Both are commonly used as part of HTTPS (HyperText Transfer Protocol Secure).

SSL was developed for the Netscape web browser in the 1990s. SSL 2.0 was the first released version; SSL 3.0 fixed a number of security issues with it. TLS was based on SSL 3.0. It is very similar to that version, with some security improvements. Although typically used for HTTPS to secure web traffic, TLS may be used for other applications such as Internet chat and email client access.

PGP

Pretty Good Privacy (PGP), created by Phil Zimmerman in 1991, brought asymmetric encryption to the masses by providing the modern suite of cryptography: confidentiality, integrity, authentication, and nonrepudiation. PGP can be used to encrypt emails, documents, or an entire disk drive. It uses a “web of trust” model to authenticate digital certificates instead of relying on a central CA.

S/MIME

MIME (Multipurpose Internet Mail Extensions) is a standard way to format email, including character sets and attachments. S/MIME (Secure MIME) leverages PKI to encrypt and authenticate MIME-encoded email. The encryption may be done by the client or the email server (called an S/MIME gateway).

Escrowed encryption

Escrowed encryption divides a private key into two or more parts. The parts are held in escrow by different, trusted third-party organizations, which only release their portion of the key with proper authorization, such as a court order.

CLIPPER CHIP

“Clipper Chip” was the name of the technology used in the Escrowed Encryption Standard (EES), an effort announced in 1993 by the U.S. government to deploy escrowed encryption in telecommunications devices. This effort created a media firestorm and was abandoned by 1996. The Clipper Chip used the Skipjack algorithm, a symmetric cipher with an 80-bit key. The algorithm was originally classified as secret.

SUMMARY OF EXAM OBJECTIVES

Cryptography provides security for data in motion and at rest. Modern systems such as Public Key Infrastructure put all of the cryptographic pieces into play via symmetric, asymmetric, and hash-based encryption to provide confidentiality, integrity, authentication, and nonrepudiation. You have learned how the pieces fit together: Slower and weaker asymmetric ciphers such as RSA and Diffie-Hellman are used to exchange faster and stronger symmetric keys such as AES and DES. The symmetric keys act as session keys to encrypt short-term sessions, such as web connections via HTTPS. Digital signatures employ public key encryption and hash algorithms such as MD5 and SHA-1 to provide nonrepudiation: sender authentication and message integrity. Understanding these concepts and others discussed in this chapter and applying them together are critical for success on the exam.

TOP FIVE TOUGHEST QUESTIONS

1. Which algorithm should you use for a low-power device that must employ digital signatures?
 - A. AES
 - B. RSA
 - C. ECC
 - D. ElGamal
2. Which of the following is true for digital signatures?
 - A. The sender encrypts the hash with a public key
 - B. The sender encrypts the hash with a private key
 - C. The sender encrypts the plaintext with a public key
 - D. The sender encrypts the plaintext with a private key
3. Which of the following attacks analyzes large amounts of plaintext/ciphertext pairs created with the same key?
 - A. Known plaintext attack
 - B. Differential cryptanalysis
 - C. Linear cryptanalysis
 - D. Chosen plaintext attack
4. What is a Hashed Message Authentication Code (HMAC)?
 - A. Encrypting a hash with a symmetric cipher
 - B. Encrypting a hash with an asymmetric cipher
 - C. A message digest
 - D. A checksum
5. Which of the following was not an AES finalist?
 - A. MARS
 - B. RC6
 - C. Serpent
 - D. Blowfish

Answers

1. Correct Answer and Explanation: C. Answer C is correct; digital signatures require asymmetric encryption. ECC is the strongest asymmetric algorithm per bit of key length. This allows shorter key lengths that require fewer CPU resources.

Incorrect Answers and Explanations: A, B, and D. Answers A, B, and D are incorrect. AES is a symmetric cipher, and symmetric ciphers are not used in digital signatures. RSA is based on factoring composite numbers into their primes, and ElGamal is based on discrete logarithms. Both methods provide roughly the same strength per bit and are far weaker per bit than ECC.
2. Correct Answer and Explanation: B. Answer B is correct; the sender generates a hash of the plaintext and encrypts the hash with a private key. The recipient decrypts the hash with a public key.

Incorrect Answers and Explanations: A, C, and D. Answers A, C, and D are incorrect. The sender encrypts the hash with the private key, not public. The plaintext is hashed, and not encrypted.

- 3.** Correct Answer and Explanation: C. Answer C is correct; linear cryptanalysis analyzes large amounts of plaintext/ciphertext pairs created with the same key, trying to deduce information about the key.
Incorrect Answers and Explanations: A, B, and D. Answers A, B, and D are incorrect. Linear cryptanalysis is a known plaintext attack, but the question references linear specifically, making A incorrect. Differential cryptanalysis seeks to find the “difference” between related encrypted plaintexts. A cryptanalyst chooses the plaintext to be encrypted during a chosen plaintext attack.
- 4.** Correct Answer and Explanation: A. Answer A is correct; a Hashed Message Authentication Code is hash-encrypted with a preshared symmetric key.
Incorrect Answers and Explanations: B, C, and D. Answers B, C, and D are incorrect. A digital signature encrypts a hash with an asymmetric cipher. A message digest is another name for a hash. A checksum is a simple hash.
- 5.** Correct Answer and Explanation: D. Answer D is correct; Blowfish was not an AES finalist (Twofish, based on Blowfish, was).
Incorrect Answers and Explanations: A, B, and C. Answers A, B, and C are incorrect. MARS, RC6, and Serpent were all AES finalists.

CHAPTER 4

Domain 4: Physical (Environmental) Security

55

Exam Objectives in this Chapter

- Perimeter defenses
- Site selection, design, and configuration
- System defenses
- Environmental controls

INTRODUCTION

Physical (environmental) security protects the Confidentiality, Integrity, and Availability of physical assets: people, buildings, systems, and data. The CISSP exam considers human safety as the most critical concern of this domain, trumping all other concerns.

PERIMETER DEFENSES

Perimeter defenses help prevent, detect, and correct unauthorized physical access. Buildings, like networks, should employ “defense-in-depth.” Any one defense may fail, so critical assets should be protected by multiple physical security controls: fences, doors, walls, locks, and so forth.

Fences

Fences may range from simple deterrents (e.g., 3-foot (1-meter) fencing) to preventive devices (e.g., an 8-foot (2.4-meter) fence topped with barbed wire). They should be designed to steer ingress and egress to controlled points such as exterior doors and gates.

Gates

Gates range in strength from ornamental, class I, designed to deter access, to class IV, designed to prevent a car from crashing through (e.g., airport and prison gates). For more information, see ASTM International’s “ASTM F2200” Standard Specification for Automated Vehicular Gate Construction (www.astm.org/Standards/F2200.htm).

Fast Facts

Here are the four classes of gate:

- Class I: Residential (home use)
- Class II: Commercial/General Access (parking garage)
- Class III: Industrial/Limited Access (loading dock for 18-wheeled trucks)
- Class IV: Restricted Access (airport, prison)

Bollards

A traffic **bollard** is a strong post designed to stop a car. The term derives from the short, strong posts (called mooring bollards) used to tie docked ships to piers.

Lights

Lights can act as both a detective and a deterrent control. They should be bright enough to illuminate the desired field of vision (the area being protected). Types of lighting include Fresnel, which is the same type originally used in lighthouses for aiming light in a specific direction.

A **lumen** is the amount of light one candle creates. Light was historically measured in foot-candles: One foot-candle is one lumen per square foot. **Lux**, based on the metric system, is more commonly used now: One lux is one lumen per square meter.

CCTV

Closed-Circuit Television (**CCTV**) is a detective device used to aid guards in detecting the presence of intruders in restricted areas. CCTVs using the normal light spectrum require sufficient visibility to illuminate the field of view visible to the camera. Infrared devices can “see in the dark” by displaying heat.

Older “tube cameras” are analog devices. Modern cameras use CCD (Charged Couple Discharge), which is digital. Cameras have mechanical irises that act similarly to human irises, controlling the amount of light entering the lens by changing the size of the aperture. Key issues include depth of field (the area in focus) and field of view (the entire area viewed by the camera). More light allows a larger depth of field because a smaller aperture places more of the image in focus. Correspondingly, a wide aperture (used in lower-light conditions) lowers the depth of field.

CCTV cameras may have other typical camera features such as pan and tilt (moving horizontally and vertically). They may display a fixed view, auto-scan (in which a given camera’s view shows for a few seconds and then another’s), or multiplexing (where multiple camera feeds are fed into one display).

Magnetic tape such as VHS is used to back up images from tube cameras. CCD cameras use DVR (Digital Video Recorder) or NVR (Network Video Recorder) for backups. NVR has the advantage of allowing centralized storage of all video data.

Exam Warning

Tube cameras are sometimes referred to as CRT (cathode ray tube). However, do not confuse CRT cameras with CRT displays: A CRT camera may be viewed on a CRT display, but they are different devices.

Locks

Locks on doors and windows are a preventive physical security control for preventing unauthorized physical access. Locks may be mechanical (e.g., key locks, combination locks) or electronic, often used with smart cards or magnetic stripe cards.

KEY LOCKS

Key locks require a physical key for unlocking. Keys may be shared or sometimes copied, which lowers the lock's accountability.

A common lock type is the pin tumbler, which has two sets of pins: driver and key. The correct key makes the pins line up with the shear line, allowing the lock tumbler (plug) to turn.

Ward, or **Warded**, **locks** must turn a key through channels (called wards); a "skeleton key" is designed to open varieties of warded locks.

In a **spring-bolt lock**, the mechanism "springs" in and out of the door jamb; the door may be closed with the spring bolt exposed. A deadbolt is rigid; the door cannot be closed when the deadbolt is unlocked. Both types extend into the strike plate in the door jamb.

COMBINATION LOCKS

Combination locks have dials that must be turned to specific numbers, in a specific order, to be opened. They are a weak form of physical access control for production environments such as data centers. Button or keypad locks also use numeric combinations. Limited accountability due to shared combinations is the primary security issue concerning these lock types.

Smart cards and magnetic stripe cards

A **smart card** is a physical access control device that is often used for electronic locks, credit card purchases, and dual-factor authentication systems. "Smart" means that the card contains a computer circuit. Another term for a smart card is *Integrated Circuit Card (ICC)*.

Smart cards may be contact or contactless. Contact cards must be inserted into a smart card reader, while contactless cards are read wirelessly. One type of contactless card technology is Radio-Frequency Identification (**RFID**). RFID cards contain tags (also called transponders) that are read by RFID transceivers.

A magnetic stripe card contains a magnetic stripe that stores information. Unlike a smart card, it is a passive device that contains no circuits.

Tailgating

In **Tailgating** (also known as *piggybacking*), an unauthorized person follows an authorized person into a building after the authorized person unlocks and opens the door.

Mantraps and turnstiles

A **mantrap** is a preventive physical control with two doors, each of which typically requires a separate form of authentication to be opened. The intruder is “trapped” between the two after entering the mantrap. Turnstiles are designed to prevent tailgating by enforcing a “one person per authentication” rule, just as they do in subway systems.

Contraband checks

Contraband checks identify objects that are prohibited to enter a secure perimeter (e.g., an airplane). Secure buildings such as on government or military installations may also employ contraband checks.

Motion detectors and other perimeter alarms

Ultrasonic and microwave motion detectors work like the Doppler radar used to predict the weather. A wave of energy is sent out, returning an “echo” when it bounces off an object. The echo is returned more quickly when a new object (such as a person walking in range of the sensor) reflects the wave.

A photoelectric motion sensor sends a beam of light across a monitored space to another photoelectric sensor. The sensor alerts when the light beam is broken.

Ultrasonic, microwave, and infrared motion sensors are active, which means that they actively send energy. A passive sensor can be thought of as “read-only.” An example is a passive infrared (PIR) sensor, which detects infrared energy created by body heat.

Doors and windows

Always consider the relative strengths and weaknesses of doors, windows, walls, floors, ceilings, and so forth. Door hinges should face inward or be otherwise protected. Externally facing hinges that are not secured pose a security risk: Attackers can remove the hinge pins with a hammer and screwdriver, allowing the door to be opened from the hinge side.

Externally facing emergency doors should be marked for emergency use only and equipped with panic bars. Use of a panic bar should trigger an alarm.

Walls, floors, and ceilings

Walls around any internal secure perimeter such as a data center should be “slab to slab,” meaning that they should start at the floor slab and run to the ceiling slab. Raised floors and drop ceilings can obscure where the walls truly start and stop. An attacker should not be able to crawl under a wall that stops at the top of the raised floor, or climb over a wall that stops at the drop ceiling.

DID YOU KNOW?

Walls should have an appropriate fire rating (the amount of time required to fail in a fire). According to the National Fire Protection Agency (NFPA) 75: Standard for the Protection of Information Technology Equipment,

The computer room shall be separated from other occupancies within the building by fire-resistant rated walls, floor, and ceiling constructed of noncombustible or limited combustible materials. The fire resistant rating shall be commensurate with the exposure, but not less than one hour.¹

Guards

Guards are a dynamic control that may be used in a variety of situations. They may aid in inspection of access credentials, monitor CCTVs, monitor environmental controls, respond to incidents, act as a deterrent (all things being equal, criminals are more likely to target an unguarded building), and much more.

Professional guards have attended advanced training and/or schooling; amateur guards (sometimes derogatively called “mall cops”) have not. The term *pseudo guard* means an unarmed security guard.

Dogs

Dogs perform perimeter defense duties, guarding a rigid “turf.” They are often used in controlled areas, such as between the exterior building wall and a perimeter fence. Dogs primarily serve as both deterrent and a detective control. Their primary drawback is legal liability.

SITE SELECTION, DESIGN, AND CONFIGURATION

Selection, Design, and Configuration describe the process of building a secure facility, such as a data center, from site selection through final design.

Site Selection Issues

Site selection is the “greenfield” process of choosing the land on which to construct a building or data center. A greenfield is an undeveloped lot that is the design equivalent of a blank canvas.

TOPOGRAPHY

Topography is the physical shape of the land: hills, valleys, trees, and the like. Highly secure sites such as on military installations will leverage (and sometimes alter) a site's topography as a defensive measure. Topography can be used to steer ingress and egress to controlled points.

UTILITY RELIABILITY

The reliability of local utilities is a critical concern in site selection. Uninterruptible Power Supplies (UPSs) provide protection against electrical failure for a short period of time. Generators provide longer protection, but require refueling to operate for extended periods of time.

CRIME

Local crime rates also factor into site selection. The primary issue is employee safety: All employees have the right to a safe working environment. Additional issues include theft of company assets.

Site design and configuration issues

Once the site has been selected, a number of design decisions must be made. Is the site to be externally marked as a data center? Is there shared tenancy in the building? Where is the telecom demarc (demarcation point)?

SITE MARKING

Many data centers are not externally marked, to avoid drawing attention to the facility (and the expensive contents within). Similar controls include attention-avoiding details such as muted building design.

SHARED TENANCY AND ADJACENT BUILDINGS

Other tenants in a building can pose security issues: They are already behind the physical security perimeter. Adjacent buildings pose a similar risk.

A crucial issue to consider in a building with shared tenancy is a shared demarc (at which the ISP's responsibility ends and the customer's begins). Most buildings have one demarc area where all external circuits enter. Access to it allows attacks on the confidentiality, integrity, and availability of all circuits and the data flowing over them.

SYSTEM DEFENSES

System defenses are one of the last lines of defense in a defense-in-depth strategy. They assume that an attacker has physical access to a device or to media containing sensitive information. These controls are the final ones protecting the data in cases where other controls may have failed.

Port controls

Modern computers can contain multiple “ports,” which may allow copying data to or from a system. Ports can be physically disabled; examples include disabling a port on a system’s motherboard, disconnecting internal wires that connect a port to the system, and physically obstructing the port itself. Ports may also be electronically locked via system policy.

Drive and tape encryption

Drive and tape encryption protect data at rest (as opposed to data in motion—that is, moving across a network), and they are one of the few controls that protect data after physical security has been breached. **Whole-disk encryption** of mobile device hard drives is recommended.

Media storage and transportation

All sensitive backup data should be stored offsite, whether transmitted via networks or physically moved as backup media. Sites using backup media should follow strict procedures for offsite rotation.

Media cleaning and destruction

All forms of media should be securely cleaned or destroyed before disposal to prevent object reuse, which is the recovery of information from previously used objects such as computer files. Object reuse attacks range from the nontechnical, such as **dumpster diving** (searching for information by rummaging through unsecured trash), to the technical, such as recovering information from unallocated blocks on a disk drive.

Paper shredders

Paper shredders cut paper to prevent object reuse. Strip-cut shredders cut into vertical strips. Cross-cut shredders are more secure, cutting both vertically and horizontally to create “confetti.”

ENVIRONMENTAL CONTROLS

Environmental controls are designed to provide a safe environment for personnel and equipment. Power, HVAC, and fire safety are considered environmental controls.

Electricity

Reliable electricity is critical for any data center, and it is one of the top priorities when selecting, building, and designing a site. Electrical faults involve short- and long-term interruption of power, as well as various cases of low and high voltage.

Crunch Time

The following are common types of electrical faults:

- *Blackout*: prolonged loss of power
- *Brownout*: prolonged low voltage
- *Fault*: short loss of power
- *Surge*: prolonged high voltage
- *Spike*: temporary high voltage
- *Sag*: temporary low voltage

SURGE PROTECTORS, UPSS, AND GENERATORS

Surge protectors shield equipment from damage due to electrical surges. They contain a circuit or fuse that is tripped during a power surge or spike, shorting the power or regulating it down to acceptable levels.

Uninterruptible Power Supplies (UPSs) provide temporary backup power in the event of a power outage. They may also “clean” the power, protecting against surges, spikes, and other forms of electrical faults.

Generators are designed to provide power for longer periods of time than UPSs can, and will run as long as fuel is available. Sufficient fuel should be stored onsite for the period the generator is expected to provide power. Refueling strategies need to consider a disaster’s effect on fuel supply and delivery.

EMI

All electricity generates magnetism, so any electrical conductor emits Electromagnetic Interference (EMI). This includes circuits, power cables, network cables, and many others. Network cables that are poorly shielded or that run too closely together may suffer crosstalk, in which magnetism from one cable “crosses” over to another, nearby cable.

Crosstalk can be mitigated via proper network cable management. Never route power cables close to network cables. The choice of network cable can also lower crosstalk: Unshielded Twisted Pair (**UTP**) cabling is far more susceptible than Shielded Twisted Pair (**STP**) or **coaxial cable**. **Fiber optic cable** uses light instead of electricity to transmit data and is not susceptible to EMI.

HVAC

HVAC (heating, ventilation, and air conditioning) keeps the air at a reasonable temperature and humidity. HVAC units operate in a closed loop, recirculating treated air. This helps reduce dust and other airborne contaminants. The units should employ positive pressure and drainage.

Data center HVAC units are designed to maintain optimum heat and humidity levels for computers. Humidity levels of 40 to 55% are recommended. A commonly recommended “set point” temperature range for a data center is 68 to 77°F (20–25°C).

STATIC AND CORROSION

Static is mitigated by maintaining proper humidity, proper grounding of all circuits, and use of anti-static sprays, wrist straps, and work surfaces. All personnel working with sensitive computer equipment such as boards, modules, or memory chips should ground themselves before performing any work.

High humidity levels can allow the water in the air to condense onto (and into) equipment, which may lead to corrosion. Both static and corrosion are mitigated by maintaining proper humidity levels.

Heat, flame, and smoke detectors

Heat detectors alert when temperature exceeds an established safe baseline. They may trigger when a specific temperature is exceeded or when the temperature changes at a specific rate.

Smoke detectors work through two primary methods: ionization and photoelectric. Ionization-based smoke detectors contain a small radioactive source that creates a small electric charge. Photoelectric sensors work in a similar fashion, except they contain an LED (Light Emitting Diode) and a photoelectric sensor that generates a small charge while receiving light. Both types alert when smoke interrupts the radioactivity or light, lowering or blocking the electric charge.

Flame detectors detect infrared or ultraviolet light emitted in fire. One drawback to this type of detection is that the detector usually requires line of site to detect the flame; smoke detectors do not have this limitation.

Evacuation routes

Evacuation routes should be prominently posted, as they are in hotel rooms. All personnel should be advised of the quickest evacuation route from their areas.

Sites should use a meeting point, where all personnel gather in the event of emergency. The two primary evacuation roles are safety warden and meeting point leader. The safety warden ensures that all personnel safely evacuate the building in the event of an emergency or drill. The meeting point leader ensures that all personnel supposed to be at the meeting point are accounted for.

Classes of fire and suppression agents

The primary safety issue in case of fire is safe evacuation. Fire suppression systems are used to extinguish fires, and different types of fires require different suppressive agents. These systems are typically designed with personnel safety as the primary concern. The following are the classes of fire:

Class A fires involve common combustibles such as wood and paper. They are the most common and should be extinguished with water or soda acid.

Class B fires are burning alcohol, oil, and other petroleum products such as gasoline. They are extinguished with gas or soda acid. You should never use water to extinguish a class B fire.

Class C fires are electrical, fed by electricity and occurring in equipment or wiring. Electrical fires are *Conductive*, and the extinguishing agent must be *non-Conductive*, such as any type of gas.

Class D fires are burning metals and are extinguished with dry powder.

Class K fires are kitchen fires involving burning oil, grease, and the like. Wet chemicals are used to extinguish them.

Table 4.1 summarizes the classes of fires.

FIRE SUPPRESSION AGENTS

All fire suppression agents work via four methods (sometimes in combination): reducing the temperature of the fire, reducing the supply of oxygen, reducing the supply of fuel, and interfering with the chemical reaction within the fire.

Water

Water suppresses fire by lowering the temperature below the kindling point (also called the ignition point). It is the safest of all suppressive agents and recommended for extinguishing common combustible fires such as burning paper or wood. It is important to cut electrical power when extinguishing a fire with water to reduce the risk of electrocution.

Soda acid

Soda acid is a mixture of acid and soda water. In addition to suppressing fire by lowering temperature, it has suppressive properties beyond plain water: It creates a foam that can float on the surface of some liquid fires, starving the oxygen supply.

Dry powder

Extinguishing a fire with dry powder (e.g., sodium chloride) works by lowering the temperature and smothering the fire, starving it of oxygen. Dry powder is primarily used to extinguish metal fires. Flammable metals include sodium, magnesium, and many others.

Table 4.1 Classes of Fire and Suppression Agents²

US	Europe	Material	Suppression agent
A	A	Ordinary combustibles such as wood and paper	Water or soda acid
B	B	Liquid	Halon/halon substitute, CO ₂ , or soda acid
B	C	Flammable gases	Halon/halon substitute, CO ₂ , or soda acid
C	E	Electrical equipment	Halon/halon substitute, CO ₂
D	D	Combustible metals	Dry powder
K	F	Kitchen (oil or fat) fires	Wet chemicals

Wet chemical

Wet chemicals are primarily used to extinguish kitchen fires (Class K in the United States; Type F in Europe), but may also be used on common combustible fires (Class A). The chemical is usually potassium acetate mixed with water. This covers a grease or oil fire in a soapy film that lowers the temperature.

CO₂

Fires require oxygen as fuel, so they may be smothered by removing the oxygen: This is how CO₂ fire suppression works. A risk associated with CO₂ is that it is odorless and colorless, and our bodies breathe it as air. By the time we begin suffocating from lack of oxygen, it is often too late. This makes CO₂ a dangerous suppressive agent, only recommended in unstaffed areas such as electrical substations.

Halon and halon substitutes

Halon extinguishes fire via a chemical reaction that consumes energy and lowers the fire's temperature. It is being phased out, and a number of replacements with similar properties are now used.

Montreal Accord Halon has ozone-depleting properties. For this reason, the 1989 Montreal Protocol (formally, the "Montreal Protocol on Substances That Deplete the Ozone Layer") banned production and consumption of new Halon in developed countries as of January 1, 1994. However, existing Halon systems may be used along with recycled Halon.

Fast Facts

Recommended replacements for Halon include the following:

- Argon
- FE-13
- FM-200
- Inergen

FE-13 is the newest of these agents and comparatively safe. It may be breathed in concentrations of up to 30%. Other Halon replacements are typically only safe up to a 10%–15% concentration

SPRINKLER SYSTEMS

Wet pipes have water up to the sprinkler heads (hence "wet"). The sprinkler head contains metal (common in older sprinklers) or a small glass bulb designed to melt or break at a specific temperature. Once that occurs, the sprinkler head opens and water flows. Each head opens independently as the trigger temperature is exceeded. Figure 4.1 shows a bulb-type sprinkler head.



FIGURE 4.1
Bulb sprinkler head.

Dry pipe systems also have closed sprinkler heads: The difference is that the pipes are filled with compressed air. The water is held back by a valve that remains closed as long as sufficient air pressure remains in the pipes. As the dry pipe sprinkler heads open, the air pressure drops in each pipe, allowing the valve to open and send water to that head.

Deluge systems are similar to dry pipes, except that the sprinkler heads are open and larger than dry pipe heads. The pipes are empty at normal air pressure; the water is held back by a deluge valve, which opens when a fire alarm (smoke, heat, or flame) triggers.

Pre-action systems are a combination of wet, dry, and deluge systems, and require two separate triggers to release water. Single interlock systems release water into the pipes when a fire alarm triggers. The water releases once the head opens. Double interlock systems use compressed air (as in dry pipes): The water will not fill the pipes until both the fire alarm triggers and the sprinkler head opens.

PORTABLE FIRE EXTINGUISHERS

All portable fire extinguishers should be marked with the type of fire they are designed to extinguish. They should be small enough to be operated by any personnel who may need to use them, which means that the old brass extinguishers are not a recommended control.

SUMMARY OF EXAM OBJECTIVES

Physical security is implicit in most security controls and is often overlooked. We must always seek balance when implementing controls from all ten domains of knowledge. All assets should be protected by multiple defense-in-depth controls that span multiple domains. For example, a file server can be protected by policy, procedures, access control, patching, antivirus, OS hardening, locks, walls, HVAC, and fire suppression (among others). A thorough and accurate risk assessment should be conducted for all assets that must be protected. Take care to ensure that no domains or controls are overlooked or neglected.

TOP FIVE TOUGHEST QUESTIONS

1. A CRT device is different from a CCD device in what way?
 - A. A CRT is an analog display; a CCD is a digital camera
 - B. A CRT is a digital display; a CCD is an analog camera
 - C. A CRT is an analog camera; a CCD is a digital camera
 - D. A CRT is a digital camera; a CCD is an analog camera

2. Which is not a valid method for detecting smoke or flame?
 - A. Temperature
 - B. Ionization
 - C. Photoelectric
 - D. Ultraviolet
3. Which type of sprinkler system would be best for an art gallery?
 - A. Wet pipe
 - B. Dry pipe
 - C. Deluge
 - D. Pre-action
4. What is the recommended agent for extinguishing a kitchen grease fire?
 - A. Dry powder
 - B. Soda acid
 - C. Wet powder
 - D. Wet chemical
5. How do dry pipe systems work?
 - A. The sprinkler heads are open; water releases when the deluge valve is opened by a fire alarm.
 - B. They release water into the pipes when a fire alarm triggers. The water releases once the sprinkler head opens.
 - C. The pipes contain water that is released when the sprinkler head opens.
 - D. The water is held back by a valve that remains closed as long as sufficient air pressure remains in the pipes. The valve opens once the sprinkler head opens and the air pressure drops in the pipes.

Answers

1. Correct Answer and Explanation: C. Answer C is correct; a Cathode Ray Tube camera is analog. A Charge Coupled Discharge camera is digital.
Incorrect Answers and Explanations: A, B, and D. Answers A, B, and D are incorrect. Answer A is a trick answer: A CRT may *also* be a display, and the question is intentionally vague on which one it is asking about (display or camera). CRT and CCD are the primary camera types, so it is implied that CRT means camera in this context. B and D are incorrect because CRTs are analog and CCD is digital.
2. Correct Answer and Explanation: A. Answer A is correct; temperature sensors are used in heat detectors, not in smoke or flame detectors.
Incorrect Answers and Explanations: B, C, and D. Answers B, C, and D are incorrect. Ionization and photoelectric sensors are used in smoke detectors. Ultraviolet sensors are used in flame detectors.
3. Correct Answer and Explanation: D. Answer D is correct; pre-action sprinkler systems lower the chance of accidental discharge by requiring two separate triggers to deploy: The sprinkler head must open and the fire alarm must trigger. These systems lower the risk of false alarms and are typically used in areas where water would cause expensive damage.
Incorrect Answers and Explanations: A, B, and C. Answers A, B, and C are incorrect. All release water after a single trigger. This increases the chance of a false alarm causing expensive damage.

4. **Correct Answer and Explanation: C.** Answer C is correct; degaussing and destroying the hard drives are considered the most secure. They offer high assurance that the data has been removed, and visual inspection of the destroyed drives provides assurance against errors made during the destruction process.
- Incorrect Answers and Explanations: **A, B, and D.** Answers A, B, and D are incorrect. They all offer weaker protection against exposure of the PII on the drives. Overwriting the disk provides reasonable protection; however, errors made during the overwriting process will not be evident from visual inspection. Deleting sensitive data simply removes the File Allocation Table (FAT) entry; the data usually remains as unallocated space. Reformatting the drives replaces the entire FAT with a new one, but the old data usually remains as unallocated space.
5. **Correct Answer and Explanation: D.** Answer D is correct; dry pipes contain compressed air and require one trigger to deploy: The sprinkler head opens. The valve opens once the air pressure drops in the pipes, releasing water.
- Incorrect Answers and Explanations: **A, B, and C.** Answers A, B, and C are incorrect. A describes deluge systems; B describes pre-action systems; C describes wet pipe systems.

Endnotes

1. www.nfpa.org/aboutthecodes/AboutTheCodes.asp?DocNum=75.
2. Classification of Portable Fire Extinguishers, URL: www.osha.gov/doc/outreachtraining/htmlfiles/extmark.html (accessed August 27, 2010).

CHAPTER 5

Domain 5: Security Architecture and Design

69

Exam Objectives in this Chapter

- Secure system design concepts
- Secure hardware architecture
- Secure operating system and software architecture
- System vulnerabilities, threats, and countermeasures
- Security models
- Evaluation methods, certification, and accreditation

INTRODUCTION

Security Architecture and Design describes fundamental logical, hardware, operating system, and software security components, and how to use them to design, architect, and evaluate secure computer systems. Understanding these fundamental issues is critical for an information security professional.

Security Architecture and Design is a three-part domain. The first part covers the hardware and software required to achieve a secure computer system. The second part covers the logical models required to keep the system secure. The third part covers evaluation models, which quantify how secure the system actually is.

SECURE SYSTEM DESIGN CONCEPTS

Secure system design transcends specific hardware and software implementations, and represents universal best practices.

Layering

Layering separates hardware and software functionality into modular tiers. The complexity of an issue such as reading a sector from a disk drive is contained within one layer (the hardware layer in this case). No one layer (such as the application layer) is directly affected by a change to any other.

Fast Facts

A generic list of security architecture layers is

1. Hardware
2. **Kernel** and device drivers
3. Operating system
4. Applications

Abstraction

Abstraction hides unnecessary details from the user. Complexity is the enemy of security: The more complex a process is, the less secure it is. That being said, computers are tremendously complex machines. Abstraction provides a way to manage complexity.

Security domains

A **security domain** is the list of objects a subject is allowed to access. More broadly defined, domains are groups of subjects and objects with similar security requirements. Confidential, secret, and top secret, for example, are three security domains used by the U.S. Department of Defense (DoD).

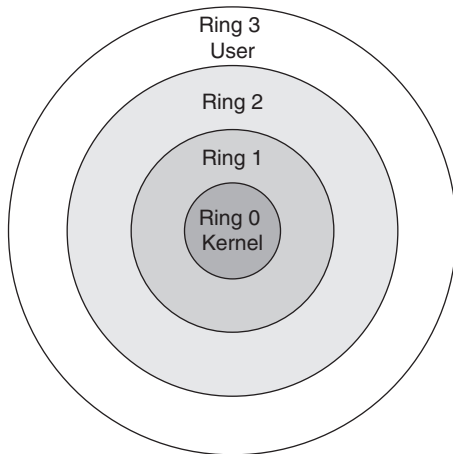


FIGURE 5.1
The ring model.

The Ring Model

The **ring model** is a form of CPU hardware layering that separates and protects domains (such as kernel mode and user mode) from each other. Many CPUs (e.g., the Intel x86 family) have four rings, ranging from ring 0 (kernel) to ring 3 (user), as shown in Figure 5.1. The innermost ring is the most trusted; each successive outer ring is less trusted.

Processes communicate between rings via system calls, which allow the processes to communicate with the kernel and provide a window between rings.

Fast Facts

The rings are (theoretically) used as follows:

- *Ring 0*: kernel
- *Ring 1*: other OS components that do not fit into Ring 0
- *Ring 2*: device drivers
- *Ring 3*: user applications

While x86 CPUs have four rings and can be used as just described, this usage is considered theoretical because most x86 operating systems, including Linux and Windows, use rings 0 and 3 only.

Open and Closed Systems

An **open system** uses open hardware and standards, employing off-the-shelf components from a variety of vendors. An IBM-compatible PC is an open system, using a standard motherboard, memory, BIOS, CPU, and so forth. You may build an IBM-compatible PC by purchasing components from a multitude of vendors. A **closed system** uses proprietary hardware or software.

SECURE HARDWARE ARCHITECTURE

Secure hardware architecture focuses on the physical computer hardware required for achieving a secure system. The hardware must provide confidentiality, integrity, and availability for processes, data, and users.

The system unit and motherboard

The **system unit** is the computer's case: It contains all of the internal electronic components: motherboard, internal disk drives, power supply, and so on. The **motherboard** contains hardware including CPU, memory slots, firmware, and peripheral slots such as PCI (Peripheral Component Interconnect) slots. The keyboard unit is the external keyboard.

The computer bus

A computer bus (shown in Figure 5.2) is the primary communication channel in a computer system. Communication between the CPU, memory, and input/output devices (e.g., keyboard, mouse, display) occurs via the bus.

The CPU

The **central processing unit (CPU)** is the "brains" of the computer, capable of controlling and performing mathematical calculations. Ultimately, everything a computer does is mathematical: adding numbers (which can be extended to subtraction, multiplication, division, and other operations), performing logical




FIGURE 5.2
Simplified computer bus.

operations, accessing memory locations by address, and the like. CPUs are rated by the number of clock cycles per second. A 2.4-GHz Pentium 4 CPU has 2.4 billion clock cycles per second.

ARITHMETIC LOGIC UNIT AND CONTROL UNIT

The **Arithmetic Logic Unit (ALU)** performs mathematical calculations; that is, it “computes.” It is fed instructions by the **control unit**, which acts as a traffic cop, sending instructions to it.

FETCH AND EXECUTE

CPUs fetch machine language instructions (e.g., “add 1 + 1”) and execute them (e.g., “add the numbers for the answer ‘2’”). “Fetch and execute” (also called “Fetch, decode, execute,” or FDX) actually takes four steps:

1. Fetch Instruction 1
2. Decode Instruction 1
3. Execute Instruction 1
4. Write (save) result 1

These four steps take one clock cycle to complete.

PIPELINING

Pipelining combines multiple steps into one combined process, allowing simultaneous fetch, decode, execute, and write steps for different instructions. Each part is called a pipeline stage; the pipeline depth is the number of simultaneous stages that may be completed at once.

Given our previous fetch and execute example of adding 1 + 1, a CPU without pipelining would have to wait an entire cycle before performing another computation. A four-stage pipeline can combine the stages of four other instructions:

1. Fetch Instruction 1
2. Fetch Instruction 2, Decode Instruction 1
3. Fetch Instruction 3, Decode Instruction 2, Execute Instruction 1
4. Fetch Instruction 4, Decode Instruction 3, Execute Instruction 2, Write (save) result 1
5. Fetch Instruction 5, Decode Instruction 4, Execute Instruction 3, Write (save) result 2
6. And so forth.

Pipelining is like an automobile assembly line: Instead of building one car at a time, from start to finish, many cars enter the assembly pipeline, and discrete phases (e.g., tire installation) occur on one car after another. This increases throughput.

INTERRUPTS

An **interrupt** indicates that an asynchronous event has occurred. CPU interrupts, a form of hardware interrupt, cause the CPU to stop processing its current task,

save the state, and begin processing a new request. When the new task is complete, the CPU finishes up the prior task.

PROCESSES AND THREADS

A **process** is an executable program and its associated data loaded and running in memory. A “heavy-weight process” (HWP) is also called a task. A parent process may spawn additional child processes called **threads**. These are light-weight processes (LWPs) that are able to share memory, resulting in lower overhead compared to HWPs.

MULTITASKING AND MULTIPROCESSING

Applications run as processes in memory and comprise executable code and data. **Multitasking** allows multiple tasks (HWPs) to run simultaneously on one CPU. Older and simpler operating systems, such as MS-DOS, are non-multitasking: They run one process at a time. Most modern operating systems, such as Linux and Windows XP, support multitasking

Exam Warning

Some sources refer to other concepts related to multitasking, including multiprogramming and multithreading. Multiprogramming is multiple programs running simultaneously on one CPU; multitasking is multiple tasks (processes) running simultaneously; multithreading is multiple threads (LWPs) running simultaneously.

Multiprogramming is an older term for *multitasking*, and many sources use the two synonymously. This book will use the term *multitasking* to refer to multiple simultaneous processes on one CPU.

Multiprocessing has a fundamental difference from multitasking: It runs multiple processes on multiple CPUs. Two versions are Symmetric Multiprocessing (SMP) and Asymmetric Multiprocessing (AMP; some sources use ASMP). In SMP systems one operating system manages all CPUs. In AMP systems there is one operating system image per CPU that essentially acts as an independent system.

CISC AND RISC

CISC (Complex Instruction Set Computer) and **RISC** (Reduced Instruction Set Computer) are two forms of CPU design. CISC uses a large set of complex machine language instructions, while RISC uses a reduced set of simpler instructions. Intel x86 CPUs (among many others) are CISC; ARM (Advanced RISC Machine), used in many cell phones and PDAs, PowerPC, Sparc, and others are RISC.

Memory

Memory is a series of on-off switches representing bits: 0s (off) and 1s (on). Memory may be chip-based or disk-based, or it may use other media such as tape.

RAM is Random Access Memory: “Random” means that the CPU may randomly access (jump to) any memory location. Sequential memory (such as tape) must sequentially read memory, beginning at offset zero, until it reaches the portion of memory desired.

Real (or primary) memory, such as RAM, is directly accessible by the CPU and is used to hold instructions and data for currently executing processes. Secondary memory (e.g., disk-based) is not directly accessible.

CACHE MEMORY

Cache memory is the fastest memory on the system, required to keep up with the CPU as it fetches and executes instructions. The data most frequently used by the CPU is stored in the cache. The fastest portion of the CPU cache is the register file, which contains multiple registers. Registers are small storage locations used by the CPU to store instructions and data.

The next fastest form of cache memory is Level 1 cache, located on the CPU itself. Level 2 cache is connected to (but outside) the CPU. Static Random Access Memory (**SRAM**) is used for cache memory.

RAM AND ROM

RAM is volatile memory used to hold instructions and data of currently running programs. It loses integrity after a loss of power. RAM memory modules are installed in slots on the computer motherboard.

Read Only Memory (**ROM**) is nonvolatile: Data stored in ROM maintains integrity after a power loss. A computer’s Basic Input Output System (**BIOS**) firmware is stored there. While ROM is read only, some ROM types may be written to via flashing, as we will see shortly in the section Flash Memory.

SRAM AND DRAM

Static Random Access Memory is expensive and fast, using small latches called “flip-flops” to store bits. Dynamic Random Access Memory (**DRAM**) stores bits in small capacitors (e.g., small batteries) and is slower and cheaper. The capacitors used by DRAM leak charge and must be continually refreshed to maintain integrity, typically every few to a few hundred milliseconds, depending on the DRAM type. Refreshing reads and writes the bits back to memory. SRAM does not require refreshing and maintains integrity as long as power is supplied.

Memory protection

Memory protection prevents one process from affecting the confidentiality, integrity, or availability of another. This is a requirement for secure multiuser systems (more than one logged-in user at a time) and multitasking systems (more than one simultaneously running process).

PROCESS ISOLATION

Process isolation is a logical control that attempts to prevent one process from interfering with another. This is a common feature among multiuser operating systems such as Linux, UNIX, or recent Microsoft Windows systems. Older operating systems such as MS-DOS provide no process isolation, which means that a crash in any MS-DOS application could crash the entire system.

HARDWARE SEGMENTATION

Hardware segmentation takes process isolation one step further by mapping processes to specific memory locations. This provides more security than (logical) process isolation alone.

VIRTUAL MEMORY

Virtual memory provides virtual address mapping between applications and hardware memory. It provides many functions, including multitasking (multiple tasks executing at once on one CPU), allowing multiple processes to access the same shared library in memory, swapping, and other functions.

Exam Warning

Virtual memory allows swapping, but it has other capabilities. In other words, virtual memory does not equal swapping.

Swapping

Swapping uses virtual memory to copy contents in primary memory (RAM) to or from secondary memory (not directly addressable by the CPU, on disk). Swap space is often a dedicated disk partition that is used to extend the amount of available memory. If the kernel attempts to access a page (a fixed-length block of memory) stored in swap space, a page fault occurs (i.e., the page is not located in RAM), and the page is “swapped” from disk to RAM.

FIRMWARE

Firmware stores small programs that do not change frequently, such as a computer’s BIOS (discussed later) or a router’s operating system and saved configuration. Various types of ROM chip may store firmware, including **PROM**, **EPROM**, and **EEPROM**.

PROM (Programmable Read Only Memory) can be written to once, typically at the factory. EPROM (Erasable Programmable Read Only Memory) and EEPROM (Electrically Erasable Programmable Read Only Memory) may be “flashed,” or erased and written to multiple times. The term “flashing” derives from the use of EPROMs: They were erased by flashing ultraviolet light on a small on-chip window.

A Programmable Logic Device (PLD) is field-programmable, which means that it is programmed after it leaves the factory. EPROMs, EEPROMS, and flash memory are examples of PLDs.

Flash memory

Flash memory (as in a USB thumb drive, for example) is a specific type of EEPROM, used for small portable disk drives. The difference is that any byte of an EEPROM may be written, while flash drives are written by (larger) sectors. This makes flash memory faster than EEPROMs, but still slower than magnetic disks.

BIOS

The IBM PC-compatible Basic Input Output System contains code in firmware that is executed when a PC is powered on. It first runs the Power-On Self-Test (POST), which performs basic tests, including verifying the integrity of the BIOS itself, testing the memory, and identifying system devices, among other tasks. Once the POST process is complete and successful, it locates the boot sector (for systems that boot off disks), which contains the machine code for the operating system kernel. The kernel then loads and executes, and the operating system boots up.

WORM STORAGE

WORM (Write Once Read Many) storage can be written to once and read many times. It is often used to support records retention for legal or regulatory compliance. WORM storage helps ensure the integrity of the data it contains: There is some assurance that it has not (and cannot be) been altered, short of destroying the media itself.

SECURE OPERATING SYSTEM AND SOFTWARE ARCHITECTURE

Secure Operating System and Software Architecture builds upon the secure hardware described in the previous section, providing a secure interface between hardware and the applications (and users) that access it. Operating systems provide memory, resource, and process management.

The kernel

The kernel, the heart of the operating system, usually runs in ring 0. It provides the interface between hardware and the rest of the operating system, including applications. As discussed previously, when an IBM-compatible PC is started or rebooted, the BIOS locates the boot sector of a storage device such as a hard drive. That boot sector contains the beginning of the software kernel machine code, which is then executed. Kernels have two basic designs: **monolithic** and **microkernel**.

A monolithic kernel is compiled into one static executable and all of it runs in supervisor mode. A microkernel is modular. It is usually smaller and has less native functionality than a typical monolithic kernel (hence “micro”), but can add functionality via loadable kernel modules.

REFERENCE MONITOR

A core function of the kernel is running the **reference monitor**, which mediates all access between subjects and objects. The kernel enforces the system’s security policy, such as preventing a normal user from writing to a restricted file (e.g., the system password file).

Virtualization

Virtualization adds a software layer between an operating system and the underlying computer hardware. This allows multiple “guest” operating systems to run simultaneously on one physical “host” computer.

Thin Clients

Thin clients are simpler than normal computer systems, which have hard drives, full operating systems, locally installed applications, and so forth. They rely on central servers, which serve applications and store associated data. Thin clients allow centralization of applications and their data, as well as the associated security costs, such as those of upgrades, patching, and data storage. Thin clients may be hardware-based (e.g., diskless workstations) or software-based (e.g., thin client applications).

DISKLESS WORKSTATIONS

A **diskless workstation** (also called diskless node) contains CPU, memory, and firmware, but no hard drive. It may include PCs, embedded devices, and others. The kernel and operating system are typically loaded via the network. Hardware UNIX X-Terminals are an example of diskless workstations.

SYSTEM VULNERABILITIES, THREATS, AND COUNTERMEASURES

System Threats, Vulnerabilities and Countermeasures describe security architecture and design vulnerabilities, and the corresponding exploits that may compromise system security. Countermeasures, or mitigating actions, which we will discuss, reduce the associated risk.

Covert channels

A **covert channel** is any communication that violates security policy. The communication channel used by malware installed on a system that locates Personally Identifiable Information (PII), such as credit card data, and sends it to a malicious server is an example. Two specific covert channel types are storage and timing.

A storage channel uses shared storage, such as a temporary directory, to allow two subjects to signal each other. A covert timing channel relies on the system clock to infer sensitive information.

The opposite of a covert channel is an overt channel: authorized communication that complies with security policy.

TOCTOU/race conditions

Time of Check/Time of Use (TOCTOU) attacks are also called race conditions: An attacker attempts to alter a condition after it has been checked by the operating system but before it is used. TOCTOU is an example of a state attack, where the attacker capitalizes on a change in operating system state.

Backdoors

A **backdoor** is a shortcut in a system which allows a user to bypass security checks (such as username/password authentication) to log in. Attackers will often install a backdoor after compromising a system.

Maintenance hooks are a type of backdoor; they are shortcuts installed by system designers and programmers to allow developers to bypass normal system checks during development (e.g., requiring users to authenticate). Maintenance hooks become a security issue if they are left in production systems.

Malicious code (Malware)

Malicious code, or **Malware**, is the generic term for any type of software that attacks an application or system. There are many types of malicious code that can cause damage to targeted systems, among them viruses, worms, trojans, and logic bombs.

Zero-day exploits are malicious code (a threat) for which there is no vendor-supplied patch (meaning that there is an unpatched vulnerability).

COMPUTER VIRUSES, WORMS, AND TROJANS

Computer viruses are malware that does not spread automatically: They require a carrier (usually human) and are frequently spread via floppy disk and (more recently) portable USB (Universal Serial Bus) memory.

Fast Facts

Types of viruses include

- *Macro*: a virus written in macro language (e.g., Microsoft Office or Microsoft Excel macros)
- *Boot sector*: a virus that infects the boot sector of a PC, which ensures that the virus loads upon system startup

Fast Facts—Cont'd

- *Stealth*: a virus that hides itself from the OS and other protective software, such as antivirus software
- *Polymorphic*: a virus that changes its signature upon infection of a new system, attempting to evade signature-based antivirus software
- *Multipartite (or multi-part)*: a virus that spreads via multiple vectors.

Worms are malware that self-propagates (i.e., spreads independently). A Trojan (also called a Trojan horse) is malware that performs two functions: one benign (such as a game) and one malicious.

ROOTKITS

A rootkit is malware that replaces portions of the kernel and/or operating system. A user-mode rootkit operates in ring 3 on most systems, replacing operating system components in “userland.”

A kernel-mode rootkit replaces the kernel, or loads malicious loadable kernel modules. Kernel-mode rootkits operate in ring 0 on most operating systems.

LOGIC BOMBS

A **logic bomb** is a malicious program that is triggered when a logical condition is met, such as after a number of transactions have been processed or on a specific date (a *time bomb* in this case). Malware such as worms often contain logic bombs, behaving in one manner and then changing tactics on a specific date and at a *specific* time.

XML

eXtensible Markup Language (XML) is designed as a standard way to encode documents and data. It is similar to, but more universal than, HTML and is used on the Web but not tied to it: XML can be used to store application configuration, to output from auditing tools, and for many other tasks. “eXtensible” means that users may use XML to define their own data formats.

Security Assertion Markup Language (SAML) is an XML-based framework for exchanging security information, including authentication data. Some forms of Single Sign-On (SSO) use SAML to exchange data.

APPLETS

Applets are small pieces of mobile code that are embedded in other software such as web browsers. Unlike HTML (HyperText Markup Language), which provides a way to display content, applets are executables.

Applets can be written in a variety of programming languages; two prominent applet languages are Java (by Oracle/Sun Microsystems) and ActiveX (by Microsoft). The term *applet* is used for Java; *control*, for ActiveX; however, they are functionally similar.

Database security

Databases present unique security challenges. The sheer amount of data that may be housed in a database requires special security consideration. As we will see shortly in the section Polyinstantiation, Inference, and Aggregation, the logic that connections database users may make by creating, viewing, and comparing records can lead to inference and aggregation attacks, requiring database security precautions such as **inference** controls and **polyinstantiation**.

POLYINSTANTIATION, INFERENCE AND AGGREGATION

Polyinstantiation allows two different objects to have the same name. Database polyinstantiation means that two rows may have the same primary key but different data.

Inference and **aggregation** occur when a user can use lower-level access to uncover restricted information. These issues occur in multiple realms, including database security.

Inference requires deduction: There is a mystery to be solved, and lower-level details provide the clues. Aggregation is a mathematical process: A user asks every question, receives every answer, and derives restricted information.

DATA MINING

Data mining searches large amounts of data to determine patterns that would otherwise be “lost in the noise.” Credit card issuers have become experts in data mining, searching millions of credit card transactions stored in their databases to discover signs of fraud. Simple data mining rules, such as “X or more purchases, in Y time, in Z places” can be used to discover credit cards that have been stolen and used fraudulently.

SECURITY MODELS

Now that we understand the logical, hardware, and software components required to have secure systems, and the risk posed to them by vulnerabilities and threats, we can move on to security models, which provide rules for secure system operation.

Bell-LaPadula model

The **Bell-LaPadula** model was originally developed for the Department of Defense. It is focused on maintaining the confidentiality of objects. Protecting confidentiality means *not* allowing users at a lower security level to access objects at a higher security level.

Lattice-based access controls

Lattice-based access control allows security controls for complex environments. For every relationship between a subject and an object, there are defined upper

Fast Facts

Bell-LaPadula includes the following rules and properties:

- *Simple Security Property*. “No read up”: a subject at a specific classification level cannot read an object at a higher classification level. Subjects with a “Secret” clearance cannot access “Top Secret” objects, for example.
- *Security Property*. “No write down”: a subject at a higher classification level cannot write to a lower classification level. For example, subjects who are logged into a Top Secret system cannot send emails to a secret system.
- *Strong Tranquility Property*. Security labels do not change while the system is operating
- *Weak Tranquility Property*. Security labels do not change in a way that conflicts with defined security properties

and lower access limits implemented by the system. This lattice, which allows reaching higher and lower data classification, depends on the need of the subject, the label of the object, and the role the subject has been assigned. Subjects have a Least Upper Bound (LUB) and Greatest Lower Bound (GLB) of access to the objects based on their lattice position. Figure 5.3 shows a lattice-based access control model. At the highest level of access is the box labeled {Alpha, Beta, Gamma}. A subject at this level has access to all objects in the lattice.

At the second tier of the lattice, we see that each object has a distinct upper and lower allowable limit. For example, if a subject has {Alpha, Gamma} access, the only viewable objects in the lattice are the Alpha and Gamma objects. Both represent the greatest lower boundary. The subject would not be able to view object Beta.

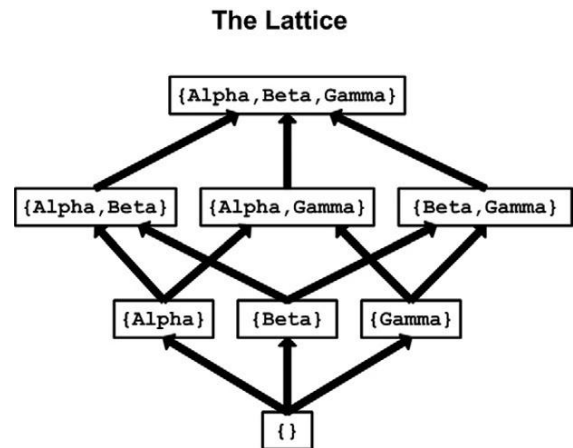


FIGURE 5.3
Lattice-based access control.

Integrity models

Models such as Bell-LaPadula focus on confidentiality, sometimes at the expense of integrity. The Bell-LaPadula “No Write Down” rule means that subjects can write up: A secret subject can write to a top secret object. What if the secret subject writes erroneous information to a top secret object? Integrity models such as Biba address this issue.

BIBA MODEL

While many governments are primarily concerned with confidentiality, most businesses want to ensure that the integrity of the information is protected at the highest level. **Biba** is the model of choice when integrity protection is vital.

Fast Facts

The Biba model has two primary rules: the Simple Integrity Axiom and the *Integrity Axiom.

- *Simple Integrity Axiom*. “No read down”: A subject at a specific classification level cannot *read* data at a lower classification. This prohibits subjects from accessing information at a lower integrity level, thus protecting integrity by preventing bad information from moving up from lower integrity levels.
- **Integrity Axiom*. “No write up”: a subject at a specific classification level cannot *write* to data at a higher classification. This prevents subjects from passing information up to an integrity level higher than the one they have clearance to change. In this way integrity is protected by preventing bad information from moving up to higher integrity levels.

Biba is often used where integrity is more important than confidentiality. Examples include time- and location-based information.

DID YOU KNOW?

Biba takes the Bell-LaPadula rules and reverses them, showing how confidentiality and integrity are often at odds. If you understand Bell-LaPadula (no read up; no write down), you can extrapolate Biba by reversing the rules (no read down; no write up).

CLARK-WILSON

Clark-Wilson is a real-world model that protects integrity by requiring subjects to access objects via programs. Because the programs have specific limits on what they can and cannot do to objects, this model effectively limits the capabilities of the subject.

Clark-Wilson uses well-formed transactions to provide integrity. This concept comprises the “access control triple”: user, transformation procedure, and constrained data item.

Chinese Wall model

The Chinese Wall model is designed to avoid conflicts of interest by prohibiting one person, such as a consultant, from accessing multiple conflict-of-interest (CoI) categories. It is also called Brewer-Nash, after its creators, Dr. David Brewer and Dr. Michael Nash, and was initially designed to address the risks inherent in employing consultants in banking and financial institutions.¹

Access control matrix

An access control matrix is a table defining the access permissions that exist between specific subjects and objects. A matrix is a data structure that acts as a table lookup for the operating system. For example, Table 5.1 is a matrix that has specific access permissions defined by users and detailing what actions they can

Table 5.1 User Access Permissions

Users	Data Access File #1	Data Creation Application
BLakey	Read/Write	Execute
AGarner	Read	Execute
CKnabe	None	None

enact. User BLakey has read/write access to the data file as well as access to the data creation application. User AGarner can read the data file and still has access to the application. User CKnabe has no access within this data access matrix.

EVALUATION METHODS, CERTIFICATION, AND ACCREDITATION

Evaluation methods and criteria are designed to gauge the real-world security of systems and products. The Trusted Computer System Evaluation Criteria (TCSEC, aka the *Orange Book*) is the oldest of the evaluation models, developed by the U.S. Department of Defense in the 1980s. Other international models have followed, including ITSEC and the Common Criteria.

The Orange Book

In 1983, the National Computer Security Center (NCSC), part of the National Institute of Standards and Technology (NIST), with help from the National Security Agency (NSA), developed the Trusted Computer System Evaluation Criteria (TCSEC). This publication is also known as the *Orange Book* because it had a bright orange cover when it was first published. TCSEC was one of the first collections of security standards implemented, and major portions of those standards are still used today in the form of U.S. Government Protection Profiles within the International Common Criteria framework.

The Trusted Network Interpretation (TNI) brings TCSEC concepts to network systems. It is often called the *Red Book* because of the color of its cover. Note that TCSEC (*Orange Book*) does not address network issues.

Begin Fast Facts

The Divisions of TCSEC include

- **D:** Minimal Protection. This division describes TCSEC-evaluated systems that do not meet the requirements of higher divisions (C through A).
- **C:** Discretionary Protection. “Discretionary” means Discretionary Access Control systems (DAC).
- **B:** Mandatory Protection. “Mandatory” means Mandatory Access Control systems (MAC).
- **A:** Verified Protection. This includes all requirements of B, plus additional controls

ITSEC

The European Information Technology Security Evaluation Criteria (ITSEC) was the first successful international evaluation model. It refers to TCSEC *Orange Book* levels, separating functionality (*F*: how well a system works) from assurance (the ability to evaluate the security of a system). There are two types of assurance: effectiveness (*Q*) and correctness (*E*).²

Assurance correctness ratings range from E0 (inadequate) to E6 (formal model of security policy). Functionality ratings include TCSEC equivalent ratings (e.g., F-C1, F-C2).

Fast Facts

The equivalent ITSEC/TCSEC ratings are

- E0: D
- F-C1, E1: C1
- F-C2, E2: C2
- F-B1, E3: B1
- F-B2, E4: B2
- F-B3, E5: B3
- F-B3, E6: A1

The International Common Criteria

The International Common Criteria represent an internationally agreed upon standard for describing and testing the security of IT products. This standard is designed to avoid requirements beyond the current state of the art; it presents a hierarchy of requirements for a range of classifications and systems. The Common Criteria are the result of the second major international information security criteria effort, following ITSEC. They use ITSEC terms such as *Target of Evaluation* and *Security Target*.

Crunch Time

The Common Criteria use specific terms when defining specific portions of the testing process:

- *Target of Evaluation* (ToE): the system or product which is being evaluated
- *Security Target* (ST): the documentation describing the TOE, including security requirements and operational environment
- *Protection Profile* (PP): an independent set of security requirements and objectives for a specific category of products or systems, such as firewalls and intrusion detection systems
- *Evaluation Assurance Level* (EAL): the evaluation score of the tested product or system

LEVELS OF EVALUATION

Within the Common Criteria are seven EALs; each builds on the level of in-depth review of the preceding level.³ For example, EAL 3-rated products can be expected to meet or exceed the requirements of products rated EAL 1 or EAL 2.

The EAL levels are described in “Common Criteria for Information Technology Security Evaluation, Part 3: Security assurance components” (July 2009, Version 3.1, Revision 3, Final, available at www.commoncriteriaportal.org/files/ccfiles/CCPART3V3.1R3.pdf).

Fast Facts

The Common Criteria levels are

- EAL 1: Functionally tested
- EAL 2: Structurally tested
- EAL 3: Methodically tested and checked
- EAL 4: Methodically designed, tested, and reviewed
- EAL 5: Semiformally designed and tested
- EAL 6: Semiformally verified, designed, and tested
- EAL 7: Formally verified, designed, and tested⁴

PCI-DSS

The Payment Card Industry Data Security Standard (**PCI-DSS**) was created by the Payment Card Industry Security Standards Council (PCI-SSC). PCI-SSC is made up of American Express, Discover, Master Card, Visa, and others. It seeks to protect credit cards by requiring vendors using them to take specific security precautions:

PCI-DSS is a multifaceted security standard that includes requirements for security management, policies, procedures, network architecture, software design, and other critical protective measures. This comprehensive standard is intended to help organizations proactively protect customer account data.⁵

SUMMARY OF EXAM OBJECTIVES

Security Architecture and Design involves the fundamental building blocks of secure computer systems, including the ring model, layers, and abstraction. We discussed secure hardware, including CPU, computer bus, RAM, and ROM. Secure software includes the kernel, reference monitor, and operating system. We use all of these together to build a secure computer system.

We learned ways to securely operate the system once built, including the Bell-LaPadula confidentiality model and the Biba integrity model, as well as modes of operation: dedicated, system high, compartmented, and multi-level secure. Finally, we learned of ways to determine assurance: proof that our systems are truly secure. Evaluation models range from TCSEC to ITSEC to the Common Criteria and beyond.

TOP FIVE TOUGHEST QUESTIONS

1. Which type of system runs multiple programs simultaneously on multiple CPUs?
 - A. Multiprocessing
 - B. Multiprogramming
 - C. Multitasking
 - D. Multithreading
2. An attacker deduces that an organization is holding an offsite meeting and so that few people are in the building, based on the low traffic volume to and from the parking lot. He uses the opportunity to break into the building and steal laptops. What type of attack has been launched?
 - A. Aggregation
 - B. Emanations
 - C. Inference
 - D. Maintenance Hook
3. What is an open system?
 - A. A process that has not been terminated
 - B. A system built from industry-standard parts
 - C. A system that allows anyone to read and change the source code
 - D. A system that contains free software
4. Which nonvolatile memory normally stores the operating system kernel on an IBM PC-compatible system?
 - A. Disk
 - B. Firmware
 - C. RAM
 - D. ROM
5. What is the highest TCSEC class applicable to a discretionary access control system that sends data across a network?
 - A. Class A
 - B. Class B
 - C. Class C
 - D. Class D

Answers

1. Correct Answer and Explanation: A. Answer A is correct; multiprocessing systems run multiple programs or processes per CPU. Two types are Symmetric Multiprocessing (SMP) and Asymmetric Multiprocessing (AMP).
Incorrect Answers and Explanations: B, C, and D. Answers B, C, and D are incorrect. All use one CPU: Multiprogramming runs multiple programs simultaneously; multitasking runs multiple tasks simultaneously; and multithreading runs multiple threads simultaneously.
2. Correct Answer and Explanation: C. Answer C is correct; inference requires an attacker to “fill in the blanks,” and deduce sensitive information from public information.

Incorrect Answers and Explanations: **A**, **B**, and **D**. Answers **A**, **B**, and **D** are incorrect. Aggregation is a mathematical operation where all questions are asked and all answers are received: There is no deduction required. Emanations are energy broadcast from electronic equipment. Maintenance hooks are system maintenance backdoors left by vendors.

- 3.** Correct Answer and Explanation: **B**. Answer **B** is correct: An open system is system hardware or software built from standard parts, such as an IBM-compatible PC.

Incorrect Answers and Explanations: **A**, **C**, and **D**. Answers **A**, **C**, and **D** are incorrect. A process that has not been terminated is running. Open Source software is source code that may be read or altered. Free software is not restricted by copyrights.

- 4.** Correct Answer and Explanation: **A**. Answer **A** is correct; the kernel is stored on disk and is loaded into volatile memory by the BIOS.

Incorrect Answers and Explanations: **B**, **C**, and **D**. Answers **B**, **C**, and **D** are incorrect. ROM (including firmware) is nonvolatile memory that stores the BIOS. RAM is volatile memory that holds the kernel after the system has booted.

- 5.** Correct Answer and Explanation: **D**. Answer **D** is correct; this is a tricky question: TCSEC (*Orange Book*) systems do not address network security, so any networked system has minimal (**D**) security. The Trusted Network Interpretation (TNI, also known as the *Red Book*) addresses network systems.

Incorrect Answers and Explanations: **A**, **B**, and **C**. Answers **A**, **B**, and **C** are incorrect. **A** (verified protection) and **B** (mandatory protection) apply to MAC (Mandatory Access Control) systems. **C** (discretionary protection) applies to DAC (Discretionary Access Control) systems with no network connection.

Endnotes

1. NIST Assessment of Access Control Systems, NIST IR 7316. URL: <http://csrc.nist.gov/publications/nistir/7316/NISTIR-7316.pdf> (accessed July 21, 2010).
2. Information Technology Security Evaluation Criteria (ITSEC)—Provisional Harmonised Criteria. URL: www.ssi.gouv.fr/site_documents/ITSEC/ITSEC-uk.pdf (accessed July 21, 2010).
3. The Common Criteria. URL: www.commoncriteriaportal.org/files/ccfiles/CCPART3V3.1R3.pdf (accessed July 21, 2010).
4. Ibid.
5. About the PCI Data Security Standard (PCI DSS). URL: www.pcisecuritystandards.org/security_standards/pci_dss.shtml (accessed July 21, 2010).

This page intentionally left blank

CHAPTER 6

Domain 6: Business Continuity and Disaster Recovery Planning

89

Exam Objectives in this Chapter

- BCP and DRP overview and process
- Developing a BCP/DRP
- DRP testing, training, and awareness
- Continued BCP/DRP maintenance
- Specific BCP/DRP frameworks

INTRODUCTION

Business Continuity Planning and Disaster Recovery Planning (BCP/DRP) together have emerged as a critical domain in the common body of knowledge. Our world of the past ten years has experienced many disruptive events: terrorism, earthquakes, hurricanes, tsunamis, floods—the list goes on.

BCP/DRP is an organization's last line of defense: When all other controls have failed, it is the final control that may prevent drastic events such as injury, loss of life, or failure of an organization. As information security professionals, we must be vigilant and protect our organizations and staff from these disruptions.

BCP AND DRP OVERVIEW AND PROCESS

The goal of this chapter is to ensure a basic understanding of the overall approach to and major phases of BCP/DRP prior to delving into the details of each phase, which we will do in the next major section, Developing a BCP/DRP.

Disasters are an inevitable fact of life. Given a long enough operational existence, every organization will experience one. A thorough, regimented, and ongoing review of the threats associated with disaster events, an organization's vulnerabilities to those threats, and the likelihood of such threats being made manifest will allow an organization to appropriately mitigate a disaster's inherent risks.

Business continuity planning

Though many organizations use the phrases “Business Continuity Planning” and “Disaster Recovery Planning” interchangeably, they are two distinct disciplines. The goal of a Business Continuity Plan (BCP) is to ensure that the business will continue to operate before, throughout, and after a disaster event. The focus of a BCP is on the business as a whole, and on ensuring that those critical services that the business provides or critical functions that it regularly performs can still be carried out both in the wake of a disruption and after the disruption has been weathered. BCP provides a long-term strategy for ensuring the continued successful operation of an organization in spite of inevitable disruptive events and disasters.

Disaster recovery planning

While Business Continuity Planning provides the long-term strategic business plan for continued operation after a disruptive event, the Disaster Recovery Plan (DRP) is more tactical in its approach, providing a short-term plan for dealing with specific IT-oriented disruptions. The DRP focuses on efficiently attempting to mitigate the impact of a disaster and the immediate response and recovery of critical IT systems in the face of significant disruption. It is considered tactical rather than strategic and provides a means for immediate response. The DRP does not focus on long-term business impact in the same fashion that a BCP does.

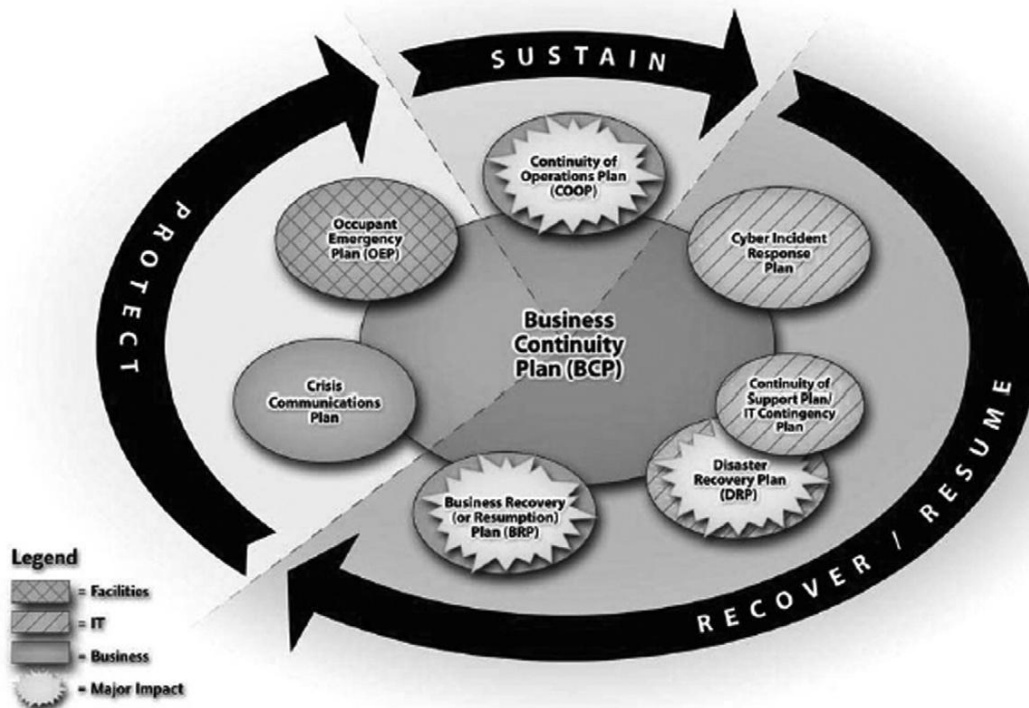
Relationship between BCP and DRP

The BCP is an umbrella for multiple specific plans, the most important is the Disaster Recovery Plan. The DRP serves as a subset of the BCP, which would be doomed to fail if it did not contain a tactical method for immediately dealing with the disruption of information systems. Figure 6.1, from NIST Special Publication 800-34, provides a visual way to understand the interrelatedness of a BCP and a DRP, as well as a Continuity of Operations Plan (COOP), an Occupant Emergency Plan (OEP), and others.

Disasters or Disruptive Events

Given that an organization’s Business Continuity and Disaster Recovery plans are created because of the potential impact of a disaster on operations, a solid understanding of the various classifications of disasters is necessary.

The analysis of threats and the determination of the associated likelihood of their being manifested are an important part of the BCP/DRP process. Table 6.1 provides a quick summary of disaster events and what type of disaster they constitute.

**FIGURE 6.1**

BCP and related plans.

Source: Swanson, M., Wohl, A., Pope, L., Grance, T., Hash, J., Thomas, R. NIST SP 800-34, *Contingency Planning Guide for Information Technology Systems*. Available from <http://csrc.nist.gov/publications/nistpubs/800-34-rev1/sp800-34-rev1.pdf> (accessed July 23, 2010).

Fast Facts

There are three common categories for the causes of disaster: natural, human, and environmental.¹

- *Natural*. Natural threats include earthquakes, hurricanes, tornadoes, floods, and some types of fires. Historically, natural disasters have been some of the most devastating that an organization must respond to.
- *Human*. Human threats represent the most common source of disasters. They can be further classified as to whether they are intentional or unintentional.
- *Environmental*. Environmental threats focus on information systems or data center environments; they are represented by power issues (blackout, brownout, surge, spike), system component or other equipment failures, and application or software flaws, among others.

Table 6.1 Examples of Disruptive Events

Disruptive event	Type
Earthquake/tornado/hurricane/etc.	Natural
Strike	Human (intentional)
Cyber terrorism	Human (intentional)/Technical
Malware	Human (intentional)/Technical
Denial of service	Human (intentional)/Technical
Errors and omissions	Human (unintentional)
Electrical fire	Environmental
Equipment failure	Environmental

Fast Facts

Types of disruptive events include

- *Errors and omissions.* Typically considered the most common source of disruptive events. This type of threat is caused by humans who are unintentionally a source of harm.
- *Natural disasters.* Earthquakes, hurricanes, floods, tsunamis, and the like.
- *Electrical or power problems.* May cause availability issues, as well as integrity issues, due to corrupted data.
- *Temperature and humidity failures.* May cause damage to equipment due to overheating, corrosion, or static electricity.
- *Warfare, terrorism, and sabotage.* Can vary dramatically based on geographic location, industry, and brand value, as well as the interrelatedness with other high-value target organizations.
- *Financially motivated attacks:* Include exfiltration of cardholder data, identity theft, pump-and-dump stock schemes, bogus anti-malware tools, corporate espionage, and others.
- *Personnel shortages.* May be caused by strikes, pandemics, or transportation issues. A lack of staff may lead to operational disruption.

The disaster recovery process

Having discussed the importance of Business Continuity and Disaster Recovery Planning as well as examples of threats that justify this degree of preparation, we will now focus on the fundamental steps involved in recovering from a disaster.

RESPOND

In order to begin the disaster recovery process, there must be an initial response—assessing the damage. Speed is essential during this initial assessment, which will determine if the event in question constitutes a disaster.

ACTIVATE TEAM

If a disaster is declared, the recovery team needs to be activated. Depending on the scope of the disaster, this communication can prove extremely difficult. Calling trees, which will be discussed later in this chapter, can

facilitate this process to ensure that team members will be activated as smoothly as possible.

COMMUNICATE

One of the most difficult aspects of disaster recovery is ensuring that consistent, timely status updates are communicated to the central team managing response and recovery. This communication often must occur out of band, meaning that the typical communication via office phones is quite often not a viable option. In addition to communication of internal status regarding recovery activities, the organization must be prepared to provide external communications, which involve disseminating details regarding the organization's recovery status with the public.

ASSESS

Though an initial assessment will have been carried out during the initial response, a more detailed and thorough assessment must be done by the now activated disaster recovery team. The team will assess the extent of the damage to determine the steps necessary to ensure the organization's ability to meet its mission.

RECONSTITUTE

The primary goal of the reconstitute phase is to successfully recover critical business operations either at a primary or a secondary site. If an alternate site is used, adequate safety and security controls must be in place in order to maintain the expected level of security the organization typically employs. In addition to the recovery team's efforts to reconstitute critical business functions at an alternate location, a salvage team will be employed to begin recovery at the primary facility, which experienced the disaster.

DEVELOPING A BCP/DRP

Developing a BCP/DRP is vital for an organization's ability to respond and recover from an interruption in normal business functions or from a catastrophic event. To ensure that all planning has been considered, the BCP/DRP has a specific set of requirements to review and implement. Listed next are the high-level steps, according to NIST 800-34, involved in achieving a sound, logical BCP/DRP. (NIST 800-34 is the National Institute of Standards and Technologies Information Technology Contingency Planning Guide, which can be found at <http://csrc.nist.gov/publications/nistpubs/800-34-rev1/sp800-34-rev1.pdf>.)²

- Project initiation
- Project scoping
- Business impact analysis
- Preventive controls identification
- Recovery strategy

- Plan design and development
- Implementation, training, and testing
- BCP/DRP maintenance

Project initiation

To develop the BCP/DRP, the scope of the project must be determined and agreed on.

Fast Facts

Project Initiation involves seven distinct milestones as quoted here³:

1. *Develop the contingency planning policy statement:* A formal department or agency policy provides the authority and guidance necessary to develop an effective contingency plan.
2. *Conduct the business impact analysis (BIA):* The BIA helps to identify and prioritize critical IT systems and components. A template for developing the BIA is also provided to assist the user.
3. *Identify preventive controls:* Measures taken to reduce the effects of system disruptions can increase system availability and reduce contingency life cycle costs.
4. *Develop recovery strategies:* Thorough recovery strategies ensure that the system may be recovered quickly and effectively following a disruption.
5. *Develop an IT contingency plan:* The contingency plan should contain detailed guidance and procedures for restoring a damaged system.
6. *Plan testing, training, and exercises:* Testing the plan identifies planning gaps, whereas training prepares recovery personnel for plan activation; both activities improve plan effectiveness and overall agency preparedness.
7. *Plan maintenance:* The plan should be a living document that is updated regularly to remain current with system enhancements.

Assess critical state

Assessing the critical state can be difficult because determining which pieces of the IT infrastructure are critical depends solely on how a piece supports users within the organization. For example, without consulting all users, a simple mapping program may not seem to be a critical asset for an organization. However, if there is a user group that makes deliveries, this mapping software may be required for scheduling them.

Conduct Business Impact Analysis

The Business Impact Analysis (BIA) is the formal method for determining how a disruption to the organization's IT system(s) will affect the organization's requirements, processes, and interdependencies with respect to the business mission.⁴ BIA identifies and prioritizes critical IT systems and components, enabling the BCP/DRP project manager to fully characterize IT contingency requirements and priorities.⁵ The objective is to correlate each IT system component with

the critical service it supports. The BIA also aims to quantify the consequence of a disruption to the component and how that will affect the organization. The primary goal is to determine the Maximum Tolerable Downtime (MTD) for a specific IT asset. This will directly impact the choice of disaster recovery solution.

Exam Warning

The BIA comprises two processes. First, identification of critical assets; second, a comprehensive risk assessment.

IDENTIFY CRITICAL ASSETS

The critical asset list contains IT assets that are deemed business-essential by the organization. These assets' DRP/BCP must have the best available recovery capabilities assigned to it.

CONDUCT BCP/DRP-FOCUSED RISK ASSESSMENT

The BCP/DRP-focused risk assessment determines which risks are inherent to which IT assets. A vulnerability analysis is also conducted for each IT system and major application because most traditional BCP/DRP evaluations focus on physical security threats, both natural and human.

DETERMINE MAXIMUM TOLERABLE DOWNTIME

The primary goal of the BIA is to determine the Maximum Tolerable Downtime (MTD), which describes the total time a system can be inoperable before the impact on the organization becomes severe. It is the maximum time it takes to complete the reconstitute phase. Reconstitution is the process of moving an organization from disaster recovery to normal business operations.

Maximum Tolerable Downtime comprises two metrics: the Recovery Time Objective (RTO) and Work Recovery Time (WRT) (to be discussed).

Alternate terms for MTD

Depending on the business continuity framework, terms that may be substituted for *Maximum Tolerable Downtime* include *Maximum Allowable Downtime* (MAD), *Maximum Tolerable Outage* (MTO), and *Maximum Acceptable Outage* (MAO).

FAILURE AND RECOVERY METRICS

A number of metrics are used to quantify how frequently systems fail, how long a system may exist in a failed state, and the maximum time to recover from failure. They include Recovery Point Objective, Recovery Time Objective, Work Recovery Time, Mean Time Between Failures, Mean Time to Repair, and Minimum Operating Requirements.

Recovery Point Objective

The Recovery Point Objective (RPO) is the level of data/work loss or system inaccessibility (measured in time) resulting from a disaster or disruptive event that an organization can withstand.

If you perform weekly backups, someone made a decision that your company could tolerate the loss of a week's worth of data. If backups are performed on Saturday evenings and a system fails on Saturday afternoon, that week's worth of data is gone. This is the recovery point objective. In this case, the RPO is one week.⁶

Recovery Time Objective and Work Recovery Time

The Recovery Time Objective (RTO) describes the maximum time allowed to recover business or IT systems. Also called systems recovery time, it is one part of Maximum Tolerable Downtime: Once the system is physically running, it must be configured.

Crunch Time

Work Recovery Time (WRT) describes the time required to configure a recovered system. “Downtime consists of two elements, the systems recovery time and the work recovery time. Therefore, $MTD = RTO + WRT$.”⁷

Mean Time between Failures

Mean Time between Failures (MTBF) quantifies how long a new or repaired system will run before failing. It is typically generated by a component vendor and is largely applicable to hardware as opposed to applications and software.

Mean Time to Repair

Mean Time to Repair (MTTR) describes how long it will take to recover a specific failed system. It is the best estimate for reconstituting the IT system to achieve business continuity.

Minimum Operating Requirements

Minimum Operating Requirements (MOR) describe the minimum environmental and connectivity requirements for computer equipment to operate. It is important to document the MOR for each IT critical asset because, in the event of a disruptive event or disaster, proper analysis can be conducted quickly to determine if the asset will be able to function in the emergency environment.

Identify preventive controls

Preventive controls prevent the potential impact of disruptive events. For example, as stated in Chapter 4, HVAC systems are designed to prevent computer equipment from overheating and failing.

DID YOU KNOW?

The BIA will identify risks that may be mitigated immediately. This is another advantage of performing BCP/DRP, including the BIA: It improves your security, even if no disaster occurs.

Recovery strategy

Once the BIA is complete, the BCP team knows the Maximum Tolerable Downtime. This metric, as well as Recovery Point Objective, Recovery Time Objective, and others, determine the recovery strategy. A cold site cannot be used if the MTD is 12 hours, for example. As a general rule, the shorter the MTD, the more expensive the recovery solution.

REDUNDANT SITE

A redundant site is an exact production duplicate of a system which has the capability to seamlessly operate all necessary IT operations without loss of services to the system's end user. It receives data backups in real time so that in the event of a disaster, users suffer no loss of data availability. A redundant site is configured exactly like the primary site and is the most expensive recovery option because it effectively more than doubles the cost of IT operations.

HOT SITE

A hot site is a data center to which an organization may relocate following a major disruption or disaster. It is equipped with a raised floor, power, utilities, computer peripherals, and fully configured computers. The hot site will have all necessary hardware and critical applications data mirrored in real time. It allows the organization to resume critical operations within a very short period of time—sometimes less than an hour.

It is important to note the difference between a hot and a redundant site. Hot sites can quickly recover critical IT functionality, perhaps even in minutes instead of hours. However, a redundant site appears to the end user as operating normally no matter what the state of operations for the IT program. A hot site has all of the physical, technical, and administrative controls implemented at the production site.

WARM SITE

A warm site has some aspects of a hot site—for example, readily accessible hardware and connectivity—but it must rely on backup data in order to reconstitute a system after a disruption. A warm site, too, is a data center

with a raised floor, power, utilities, computer peripherals, and fully configured computers.

COLD SITE

A cold site is the least expensive recovery solution. It does not include backup copies of data nor does it contain any immediately available hardware. After a disruptive event, a cold site takes the longest time of all recovery solutions to be implemented and to restore critical IT services. Especially in a disaster area, it may take weeks to obtain and install vendor hardware. Organizations using this recovery solution will thus have to be able to withstand a significantly long MTD—usually measured in weeks, not days. A typical cold site data center has a raised floor, power, utilities, and physical security, but not much beyond that.

RECIPROCAL AGREEMENT

Reciprocal agreements are bi-directional arrangements between two organizations in which one promises the other that the latter can move in and share space in the wake of a disaster. Such agreements are documented in the form of a contract written to gain support in a disaster from outside organizations. They are also referred to as Mutual Aid Agreements (MAAs) and are structured so that the parties will assist the other in an emergency.

MOBILE SITE

Mobile sites are “data centers on wheels”—that is, towable trailers that contain racks of computer equipment as well as HVAC, fire suppression, and physical security. They are a good fit for disasters such as a flood, where the data center is damaged but the rest of the facility and surrounding property are intact. They may be towed onsite, supplied power and network, and brought online.

Related plans

As discussed previously, the Business Continuity Plan is an umbrella that covers other plans. In addition to the disaster recovery plan, these include

- Continuity of Operations Plan (COOP)
- Business Resumption/Recovery Plan (BRP)
- Continuity of Support Plan
- Cyber Incident Response Plan
- Occupant Emergency Plan (OEP)
- Crisis Management Plan (CMP)

Table 6.2, from NIST Special Publication 800-34,⁸ provides a summary.

Table 6.2 Summary of BCP Plans

Plan	Purpose	Scope
Business Continuity Plan (BCP)	Provide procedures for sustaining essential business operations while recovering from a significant disruption	Addresses business processes; IT addressed based only on its support for business process
Business Recovery (or Resumption) Plan (BRP)	Provide procedures for recovering business operations immediately following a disaster	Addresses business processes; not IT-focused; IT addressed based only on its support for business process
Continuity of Operations Plan (COOP)	Provide procedures and capabilities to sustain an organization's essential, strategic functions at an alternate site for up to 30 days	Addresses the subset of an organization's missions that are deemed most critical; usually written at headquarters level; not IT-focused
Continuity of Support Plan/IT Contingency Plan	Provide procedures and capabilities for recovering a major application or general support system	Same as IT contingency plan; addresses IT system disruptions; not business process focused
Crisis Communications Plan	Provides procedures for disseminating status reports to personnel and the public	Addresses communications with personnel and the public; not IT-focused
Cyber Incident Response Plan	Provide strategies to detect, respond to, and limit consequences of a malicious cyber incident	Focuses on information security responses to incidents affecting systems and/or networks
Disaster Recovery Plan (DRP)	Provide detailed procedures to facilitate recovery of capabilities at an alternate site	Often IT-focused; limited to major disruptions with long-term effects
Occupant Emergency Plan (OEP)	Provide coordinated procedures for minimizing loss of life or injury and protecting property damage in response to a physical threat	Focuses on personnel and property particular to the specific facility; not business process or IT system functionality based

Call trees

A key tool for staff communication in the Crisis Communications Plan is the call tree, which is used to quickly communicate news throughout an organization without overburdening any one employee. The call tree works by assigning each employee a small number of other employees whom they must call in an emergency. For example, the president may notify his board of directors of an emergency situation and they, in turn, notify their top-tier managers. The top-tier managers then notify the subordinates they have been assigned. The call tree continues until all affected personnel have been contacted.

DRP TESTING, TRAINING, AND AWARENESS

Testing, training, and awareness must be in place during the “disaster” portion of a BCP/DRP. Skipping them is one of the most common BCP/DRP mistakes.

DRP testing

To ensure that a Disaster Recovery Plan is viable, thorough testing is needed. Further, given the DRP’s detailed tactical subject matter, it should come as no surprise that routine infrastructure, hardware, software, and configuration changes materially alter the way in which the DRP will be carried out. Although an organization’s information systems are in a constant state of flux, unfortunately not many changes readily make their way into an updated DRP. To ensure both the initial and continued efficacy of the DRP as a recovery methodology, testing needs to be performed.

REVIEW

The DRP Review is the most basic initial DRP test, focusing on a reading of the DRP in its entirety to ensure complete coverage. This review is typically performed by the team that developed the plan and involves team members reading the entire plan quickly to uncover any obvious flaws. The DRP Review is primarily a sanity check to ensure that there are no glaring omissions in coverage or fundamental shortcomings in approach.

CHECKLIST

Checklist (also known as *consistency*) testing lists all necessary components required for successful recovery to ensure that they are, or will be, available should a disaster occur. For example, if the DRP calls for reconstituting systems from tape backups at an alternate computing facility, does the site have an adequate number of tape drives on hand to carry out the recovery in the indicated time? The checklist test is often performed concurrently with the structured walkthrough or with tabletop testing as a solid first threshold. It is focused on ensuring that the organization has, or can acquire in a timely fashion, sufficient resources for successful recovery.

STRUCTURED WALKTHROUGH/TABLETOP

Another test that is commonly completed at the same time as the checklist test is the structured walkthrough, which is often also referred to as a tabletop exercise. In the the structured walkthrough, usually performed prior to more in-depth testing, the goal is to allow individuals knowledgeable about the systems and services targeted for recovery to thoroughly review the overall approach. The term *structured walkthrough* is illustrative: The group will walk through the proposed recovery procedures in a structured manner to determine whether there are any noticeable omissions, gaps, erroneous assumptions, or simply technical missteps that would hinder the recovery process.

SIMULATION TEST/WALKTHROUGH DRILL

A simulation test, also called a walkthrough drill (not to be confused with the discussion-based structured walkthrough) goes beyond talking and actually has teams carry out the recovery process. A disaster is simulated, to which the team must respond as they are directed by the DRP. The scope of simulation varies significantly and tends to grow more complicated, and involve more systems, as smaller disaster simulations are successfully managed. Though some will see the goal as being successful recovery of systems impacted by the simulation, ultimately the goal of any DRP test is to ensure that the organization is well prepared in the event of an actual disaster.

PARALLEL PROCESSING

Another type of DRP test, parallel processing, is common in environments where transactional data is a key component of processes critical to the business. Typically it will involve recovery of crucial processing components at an alternate computing facility, and then restoration of data from a previous backup. Note that regular production systems are not interrupted.

PARTIAL AND COMPLETE BUSINESS INTERRUPTION

Arguably the most high fidelity of all DRP tests involves business interruption. However, this test can actually be the cause of a disaster, so extreme caution should be exercised before attempting it. As the name implies, in business interruption testing the organization actually stops processing normal business at the primary location and instead uses the alternate computing facility. Such tests are more common in organizations where fully redundant, often load-balanced, operations already exist.

Training

Although an element of DRP training comes with performing the tests just discussed, there is certainly a need for more detailed training on specific elements of the DRP process. Another reason for training is to ensure adequate representation of staff trained in basic first aid and CPR.

STARTING EMERGENCY POWER

Though it might seem simple, converting a data center to emergency power, such as backup generators that take the load as the UPS fails, is not to be taken lightly. Specific training for and testing of failing over to emergency power should be regularly performed.

CALL TREE TRAINING/TESTING

Another example of combination training and testing is call trees, which were discussed previously. The hierarchical relationships of call trees can make outages in the tree problematic. Individuals with calling responsibilities are expected to be able to answer within a very short time or make other arrangements.

CONTINUED BCP/DRP MAINTENANCE

Once the initial BCP/DRP plan is completed, tested, trained, and implemented, it must be kept up to date. Business and IT systems change quickly, and IT professionals are accustomed to adapting. BCP/DRP plans must keep pace with all critical business and IT changes.

Change management

The change management process, discussed in depth in Chapter 9, is designed to ensure that security is not adversely affected as systems are introduced, modified, and updated. Change Management includes tracking and documenting all planned changes, formal approval of substantial changes, and documentation of the results of completed changes. All changes must be auditable.

Crunch Time

The BCP team should be a member of the change control board and attend all meetings. The goal of the BCP

team's involvement is to identify any changes that must be addressed by the BCP/DRP plan.

BCP/DRP mistakes

Business continuity and disaster recovery planning is an organization's last line of defense against failure. If other controls have failed, BCP/DRP is the final recourse. If it fails, the business may fail.

The success of BCP/DRP is critical, but many do plans fail. The BCP team should consider the failure of other organizations' plans, and place their own under intense scrutiny. They should ask themselves this question: "Have we made mistakes that threaten the success of our plan?"

Fast Facts

Common BCP/DRP mistakes include

- Lack of management support
- Lack of business unit involvement
- Lack of prioritization among critical staff
- Improper (often overly narrow) scope
- Inadequate telecommunications management
- Inadequate supply chain management
- Incomplete or inadequate crisis management plan
- Lack of testing
- Lack of training and awareness
- Failure to keep the BCP/DRP plan up to date

SPECIFIC BCP/DRP FRAMEWORKS

Given the patchwork of overlapping terms and processes used by various BCP/DRP frameworks, this chapter focuses on universal best practices, without attempting to map to a number of different (and sometimes inconsistent) terms and processes that these frameworks describe.

A handful of specific frameworks are worth discussing, including NIST SP 800-34, ISO/IEC-27031, BCI, and BS-25999.

NIST SP 800-34

The National Institute of Standards and Technology (NIST) Special Publication 800-34 Contingency Planning Guide for Information Technology Systems may be downloaded at <http://csrc.nist.gov/publications/nistpubs/800-34-rev1/sp800-34-rev1.pdf>. The document is high quality and in the public domain. Plans can sometimes be significantly improved by referencing SP 800-34 when they are written or updated.

ISO/IEC-27031

ISO/IEC-27031 is a draft guideline that is part of the ISO 27000 series, which also includes ISO 27001 and ISO 27002 (discussed in Chapter 1). Because it is still in draft form, this document should not (yet) be used as a primary resource for studying BCP/DRP for the exam. This may change, however, once ISO/IEC-27031 becomes a full standard. It focuses on BCP (DRP is handled by another framework that is discussed later).

The current formal name is ISO/IEC 27031 Information technology—Security Techniques—Guidelines for ICT Readiness for Business Continuity (final committee draft).

Fast Facts

According to www.iso27001security.com/html/27031.html site, as quoted here, ISO/IEC 27031 is designed to

- Provide a framework (methods and processes) for any organization—private, governmental, and non-governmental
- Identify and specify all relevant aspects including performance criteria, design, and implementation details, for improving ICT readiness as part of the organization's ISMS, helping to ensure business continuity.
- Enable an organization to measure its continuity, security and hence readiness to survive a disaster in a consistent and recognized manner.⁹

Terms and acronyms used by ISO/IEC 27031 include

- ICT—Information and Communications Technology
- ISMS—Information Security Management System

A separate ISO plan for disaster recovery is “ISO/IEC 24762:2008, Information Technology—Security techniques—Guidelines for Information and Communications Technology Disaster Recovery Services.” More information is available at www.iso.org/iso/catalogue_detail.htm?csnumber=41532.

BS-25999

The British Standards Institution (BSI; www.bsigroup.co.uk/) released BS-25999, which is in two parts:

- Part 1, the Code of Practice, provides business continuity management best practice recommendations. Please note that this is a guidance document only.
- Part 2, the Specification, provides the requirements for a Business Continuity Management System (BCMS) based on BCM best practice. This is the part of the standard that you can use to demonstrate compliance via an auditing and certification process.¹⁰

BCI

The Business Continuity Institute (BCI, www.thebci.org/) published its six-step Good Practice Guidelines (GPG) in 2008, describing the Business Continuity Management (BCM) process:

- Section 1 consists of the introductory information plus BCM Policy and Programme Management.
- Section 2 is Understanding the Organisation
- Section 3 is Determining BCM Strategy
- Section 4 is Developing and Implementing BCM Response
- Section 5 is Exercising, Maintaining & Reviewing BCM arrangements
- Section 6 is Embedding BCM in the Organisation’s Culture¹¹

The 2008 GPG cross-references BS25999. The 2010 GPG states “There are no longer any cross references to BS25999 and no implied direct correlation between GPG2010 and BS25999.”¹²

SUMMARY OF EXAM OBJECTIVES

Business Continuity and Disaster Recovery Planning is a critical, and frequently overlooked, domain, yet it can be the most critical of all, and can serve as an organization’s last control to prevent failure. Of all controls, a failed BCP or DRP can be the most devastating, potentially resulting in organizational failure or injury or loss of life.

Beyond mitigating such stark risks, Business Continuity and Disaster Recovery Planning has evolved to provide true business value to organizations, even absent a disaster. The organizational diligence required to build a comprehensive BCP/

DRP can pay many dividends through a thorough understanding of key business processes, asset tracking, prudent backup and recovery strategies, and the use of standards. Mapping risk to key business processes can result in preventive risk measures taken in advance of any disaster, which may avoid future disasters entirely.

TOP FIVE TOUGHEST QUESTIONS

1. Which plan details the steps required to restore normal business operations after recovery from a disruptive event?
 - A. Business Continuity Plan (BCP)
 - B. Business Resumption Plan (BRP)
 - C. Continuity of Operations Plan (COOP)
 - D. Occupant Emergency Plan (OEP)
2. Which metric describes how long it will take to recover a failed system?
 - A. Minimum Operating Requirements (MOR)
 - B. Mean Time between Failures (MTBF)
 - C. Mean Time to Repair (MTTR)
 - D. Recovery Point Objective (RPO)
3. Which metric describes the moment in time at which data must be recovered and made available to users in order to resume business operations?
 - A. Mean Time between Failures (MTBF)
 - B. Mean Time to Repair (MTTR)
 - C. Recovery Point Objective (RPO)
 - D. Recovery Time Objective (RTO)
4. Maximum Tolerable Downtime (MTD) comprises which two metrics?
 - A. Recovery Point Objective (RPO) and Work Recovery Time (WRT)
 - B. Recovery Point Objective (RPO) and Mean Time to Repair (MTTR)
 - C. Recovery Time Objective (RTO) and Work Recovery Time (WRT)
 - D. Recovery Time Objective (RTO) and Mean Time to Repair (MTTR)
5. Which draft Business Continuity guideline ensures continuity of Information and Communications Technology (ICT) as part of the organization's Information Security Management System (ISMS)?
 - A. BCI
 - B. BS-7799
 - C. ISO/IEC-27031
 - D. NIST Special Publication 800-34

Answers

1. Correct Answer and Explanation: **B.** Answer **B** is correct; Business Resumption Planning details the steps required to restore normal business operations after recovery from a disruptive event.
Incorrect Answers and Explanations: **A, C,** and **D.** Answers **A, C,** and **D** are incorrect. Business Continuity Planning develops a long-term plan to ensure the continuity of business operations. The Continuity of Operations Plan describes the procedures required to maintain operations during a

disaster. The Occupant Emergency Plan provides the response procedures for occupants of a facility in the event a situation poses a threat to the health and safety of personnel, the environment, or property.

2. Correct Answer and Explanation: C. Answer C is correct; Mean Time to Repair (MTTR) describes how long it will take to recover a failed system. It is the best estimate for reconstituting the IT system to achieve business continuity.

Incorrect Answers and Explanations: A, B, and D. Answers A, B, and D are incorrect. Minimum Operating Requirements describe the minimum environmental and connectivity requirements for operating computer equipment. Mean Time between Failures quantifies how long a new or repaired system will run before failing. Recovery Point Objective is the moment in time at which data must be recovered and made available to users in order to resume business operations.

3. Correct Answer and Explanation: C. Answer C is correct; Recovery Point Objective is the moment in time at which data must be recovered and made available to users in order to resume business operations.

Incorrect Answers and Explanations: A, B, and D. Answers A, B, and D are incorrect. Mean Time Between Failures quantifies how long a new or repaired system will run before failing. Mean Time to Repair describes how long it will take to recover a failed system. Recovery Time Objective describes the maximum time allowed to recover business or IT systems.

4. Correct Answer and Explanation: C. Answer C is correct; Recovery Time Objective, the time it takes to bring a failed system back online, and Work Recovery Time, the time required to reconfigure a failed system, are used to calculate Maximum Tolerable Downtime: $RTO + WRT = MTD$.

Incorrect Answers and Explanations: A, B, and D. Answers A, B, and D are incorrect. Maximum Tolerable Downtime does not directly use Recovery Point Objective or Mean Time to Repair as a metric.

5. Correct Answer and Explanation: C. Answer C is correct; The ISO/IEC-27031 guideline ensures continuity of the Information and Communications Technology as part of the organization's Information Security Management System.

Incorrect Answers and Explanations: A, B, and D. Answers A, B, and D are incorrect. BCI and NIST Special Publication 800-34 are business continuity frameworks, but they do not match the terms in the question. BB-7799 is not BCP/DRP focused: It describes information security management best practices.

Endnotes

1. Swanson, M., Wohl, A., Pope, L., Grance, T., Hash, J., Thomas, R. NIST SP 800-34 Contingency Planning Guide for Information Technology Systems. URL: <http://csrc.nist.gov/publications/nistpubs/800-34-rev1/sp800-34-rev1.pdf> (accessed July 23, 2010).
2. Ibid.
3. Ibid.
4. Ibid.

5. Ibid.
6. Snedaker, Susan. Understanding Security Risk Management: Recovery Time Requirements. URL: http://searchsecuritychannel.techtarget.com/generic/0,295582,sid97_gci1268749,00.html (accessed July 23, 2010).
7. Ibid.
8. NIST SP 800-34.
9. ISO/IEC 27031 Information Technology—Security Techniques—Guidelines for ICT Readiness for Business Continuity (final committee draft). URL: www.iso27001security.com/html/27031.html (accessed July 23, 2010).
10. BS 25999 Business continuity. URL: www.bsigroup.co.uk/en/Assessment-and-Certification-services/Management-systems/Standards-and-Schemes/BS-25999/ (accessed July 23, 2010).
11. Business Continuity Institute. Business Continuity Management Good Practice Guidelines 2008. URL: www.thebci.org/gpg/GPG2008-2Section6FINAL.pdf (accessed July 23, 2010).
12. Business Continuity Institute. BCI Good Practice Guidelines 2010. URL: www.thebci.org/gpg.htm (accessed July 23, 2010).

This page intentionally left blank

CHAPTER 7

Domain 7: Telecommunications and Network Security

109

Exam Objectives in this Chapter

- Network architecture and design
- Network devices
- Secure communications

INTRODUCTION

Telecommunications and Network Security are fundamental to modern life. The Internet, the Web, online banking, instant messaging, email, and many other technologies rely on it: Our world could not exist without it. Telecommunications and Network Security (often called “telecommunications,” for short) is one of the largest domains in the Common Body of Knowledge, and it contains more concepts than does any other.

NETWORK ARCHITECTURE AND DESIGN

In this section we discuss how networks should be designed and the controls they may contain, focusing on deploying defense-in-depth strategies and weighing the cost and complexity of a network control against the benefit provided.

Fundamental network concepts

Before we can discuss specific Telecommunications and Network Security concepts, we need to understand the fundamentals behind them. Terms like “broadband” are often used informally: The exam requires a precise understanding of information security terminology.

SIMPLEX, HALF-DUPLEX, AND FULL-DUPLEX COMMUNICATION

Simplex communication is one way, similarly to a car radio tuned to a music station. Half-duplex communication sends or receives at one time only (not simultaneously), similarly to a walkie-talkie. Full-duplex communication sends and receives simultaneously, similarly to two people having a face-to-face conversation.

BASEBAND AND BROADBAND

Baseband networks have one channel and can send only one signal at a time. Ethernet networks are baseband: A “100baseT” **UTP** (Unshielded Twisted Pair) cable means 100-megabit, baseband, and twisted pair. Broadband networks have multiple channels and can send multiple signals at one time (e.g., cable TV). The term “channel” derives from older communications formats such as radio.

LANS, MANS, WANS, GANS, AND PANS

A **LAN** is a Local Area Network. It is comparatively small, typically confined to a building or to an area within one. A **MAN** is a Metropolitan Area Network and is typically confined to a city, a zip code, or a campus or office park. A **WAN** is a Wide Area Network, typically covering cities, states, or countries. A **GAN** is a Global Area Network—that is, a global collection of WANs.

At the other end of the spectrum, the smallest of these networks are **PANs**: Personal Area Networks, with a range of 100 meters or much smaller. Low-power wireless technologies such as Bluetooth use PANs.

INTERNET, INTRANET, AND EXTRANET

The Internet is a global collection of peered networks running TCP/IP, providing best effort service. An Intranet is a private, usually company-owned network running TCP/IP. An Extranet is a connection, such as business-to-business, between private Intranets.

CIRCUIT-SWITCHED AND PACKET-SWITCHED NETWORKS

The original voice networks were circuit switched: A dedicated circuit or channel (a portion of a circuit) was dedicated between two nodes. Circuit-switched networks can provide dedicated bandwidth to point-to-point connections, such as a T1 connecting two offices.

Instead of using dedicated circuits, Packet-switched networks divide data into packets, each sent individually. If multiple routes are available between two points on a network, packet switching can choose the best one and fall back to secondary routes in case of failure. Packets may take any path (and different paths) across a network and are reassembled by the receiving node. Missing packets can be retransmitted, and out of-order packets can be resequenced.

The OSI Model

The Open System Interconnection (OSI) Reference Model is a layered network model that was developed by the International Organization for Standardization. It is abstract—that is, we do not run it directly in our systems (most now use the TCP/IP model); rather, we use it as a reference point, so “Layer 1” (physical) is universally understood, whether, for example, we are running Ethernet or Asynchronous Transfer Mode (ATM).

The OSI model has seven layers, as shown in Table 7.1. The layers may be listed in top-to-bottom or bottom-to-top order.

Table 7.1 The OSI Model

7	Application
6	Presentation
5	Session
4	Transport
3	Network
2	Data Link
1	Physical

LAYER 1—PHYSICAL

The physical layer of the OSI model describes units of data such as **bits** represented by energy (e.g., light, electricity, radio waves) and the medium used to carry them (e.g., copper, fiber optic cables). WLANs have a physical layer, even though we cannot actually touch it. Layer 1 devices include hubs and repeaters.

LAYER 2—DATA LINK

The data link layer handles access to the physical layer as well as local area network communication. An **Ethernet** card and its Media Access Control (**MAC**) address are at Layer 2, as are switches and bridges.

LAYER 3—NETWORK

The network layer describes routing: moving data from a system on one LAN to a system on another. IP addresses and routers exist here. Layer 3 protocols include **IPv4** and **IPv6**, among others.

LAYER 4—TRANSPORT

The transport layer handles packet sequencing, flow control, and error detection. TCP and User Datagram Protocol (**UDP**) are Layer 4 protocols.

Layer 4 makes a number of features available, such as resending or resequencing packets. Using these features is a protocol implementation decision. **TCP** takes advantage of them at the expense of speed. Many of the Layer 4 features are not implemented in **UDP**, which chooses speed over reliability.

LAYER 5—SESSION

The session layer manages sessions, which provide maintenance on connections. Mounting a file share via a network requires a number of maintenance sessions, such as Remote Procedure Calls (**RPCs**): These reside at the session layer. A good way to remember the session layer's function is "connections between applications." Layer 5 uses simplex, half-duplex, and full-duplex communication.

LAYER 6—PRESENTATION

The presentation layer presents data to the application (and user) in a comprehensible way. Its concepts include data conversion; characters sets such as ASCII; and image formats such as GIF (Graphics Interchange Format), JPEG (Joint Photographic Experts Group), and TIFF (Tagged Image File Format).

LAYER 7—APPLICATION

The application layer is where the user interfaces with a computer application. The web browser, word processor, and instant messaging client exist at Layer 7. The protocols Telnet and FTP are application layer protocols.

The TCP/IP model

TCP/IP (Transmission Control Protocol/Internet Protocol) is a popular network model that was created by the U.S. Defense Advanced Research Projects Agency in the 1970s (see www.isoc.org/internet/history/brief.shtml for more information). It is simpler than the OSI model, as shown in Table 7.2.

Although TCP and IP receive top billing, TCP/IP as a whole is actually a suite of protocols including UDP and ICMP (Internet Control Message Protocol), among many others.

NETWORK ACCESS LAYER

TCP/IP's network access layer combines layers 1 (physical) and 2 (data link) of the OSI model. It addresses Layer 1 issues such as energy, bits, and the medium used to carry them (copper, fiber, wireless, etc). It also addresses Layer 2 issues such as converting bits into protocol units—for example, Ethernet frames, MAC addresses, and Network Interface Cards (NICs).

INTERNET LAYER

The Internet layer of the TCP/IP model aligns with Layer 3 (network) of the OSI model. This is where IP addresses and routing reside. When data is transmitted

Table 7.2 OSI versus TCP/IP

OSI Model		TCP/IP Model
7	Application	Application
6	Presentation	
5	Session	Host-to-Host Transport
4	Transport	
3	Network	Internet
2	Data Link	Network Access
1	Physical	

from a node on one LAN to a node on another one, the Internet layer is used. IPv4 and IPv6 (among others) are Internet Layer TCP/IP protocols.

Exam Warning

OSI's Layer 3 is called "Network." Do not confuse it with the network access TCP/IP layer, which aligns with layers 1 and 2 of the OSI model.

HOST-TO-HOST TRANSPORT LAYER

The host-to-host transport layer (sometimes called "host-to-host" or, more commonly, just "transport," which we use here) connects the Internet layer to the application layer. It is where applications are addressed on a network via ports. TCP and UDP are TCP/IP's two transport layers.

APPLICATION LAYER

The TCP/IP application layer combines OSI layers 5 through 7 (session, presentation, and application). Most of these protocols use a client-server architecture, where a client (such as `ssh`) connects to a listening server (called a daemon on UNIX systems) such as `sshd`. The clients and servers use either TCP or UDP (and sometimes both) as a transport layer protocol. TCP/IP application layer protocols include, among many others, `SSH`, `Telnet`, and `FTP`.

Network access, internet, and transport layer protocols

TCP/IP is a protocol suite: IPv4 and IPv6 at Layer 3, TCP and UDP at Layer 4, and a multitude of protocols at layers 5 through 7, including Telnet, FTP, SSH, and many others. We will focus on the lower-layer protocols, spanning the network access and transport layers.

MAC ADDRESSES

A Media Access Control address is the unique hardware address of an Ethernet Network Interface Card, typically "burned in" at the factory. MAC addresses may be changed in software.

DID YOU KNOW?

Historically, MAC addresses were 48 bits long, divided in halves: The first 24 bits represented the Organizationally Unique Identifier (**OUI**); the last 24 bits, a serial number (formally called an extension identifier). Newer EUI-64 MAC addresses are 64 bits long.

IPv4

IPv4 is Internet Protocol version 4, commonly called "IP." It is simple, designed to carry data across networks. It is also connectionless and unreliable, providing

“best effort” packet delivery. If connections or reliability are required, they must be provided by a higher-level protocol carried by IP, such as TCP.

IPv4 uses 32-bit source and destination addresses, usually shown in “dotted quad” format (e.g., “192.168.2.4”). A 32-bit address field allows 2^{32} , or nearly 4.3 billion, addresses. Too few available IPv4 addresses in a world where humans (and their devices) outnumber them is a fundamental problem: This was one of the factors leading to the creation of IPv6, which uses much larger, 128-bit, addresses.

IPv6

IPv6 is the successor to IPv4, featuring a larger address space, simpler routing, and simpler address assignment. It has become more prevalent since the release of both the Microsoft Vista and Windows 7 operating systems, which support it and have it enabled by default. Most modern Linux operating systems, such as Ubuntu, also enable IPv6 by default.

DID YOU KNOW?

Systems may be “dual stack” and use both IPv4 and IPv6 simultaneously. Hosts may also access IPv6 networks via IPv4; this is called tunneling.

TCP

TCP is the Transmission Control Protocol, a reliable Layer 4 protocol that uses a three-way handshake to create reliable connections across a network. TCP can reorder segments that arrive out of order and retransmit missing segments.

TCP Ports

TCP connects from a source port to a destination port—for example, source port 51178 and destination port 22. The TCP port field is 16 bits, allowing port numbers from 0 to 65535.

The two types of ports are reserved and ephemeral. Reserved ports are 1023 or lower; ephemeral ports are 1024 through 65535. Most operating systems require super-user privileges to open a reserved port, but any user may open an (unused) ephemeral port.

UDP

UDP is the User Datagram Protocol, a simpler and faster cousin of TCP. It is commonly used for applications that are “lossy” (i.e., they can handle some packet loss), such as streaming audio and video. It is also used for query-response applications, such as DNS queries.

ICMP

ICMP is the Internet Control Message Protocol, a helper protocol that assists Layer 3 (IP) by troubleshooting and reporting error conditions: Without ICMP, IP would fail when faced with problems like routing loops, ports, hosts, or downed

networks. ICMP has no concept of ports, as TCP and UDP do, but instead uses types and codes. Commonly used ICMP types are echo request and echo reply (used for ping) and time to live exceeded in transit (used for traceroute).

Application-layer TCP/IP protocols

A multitude of protocols exist at TCP/IP's application layer, which combines the presentation, session, and application layers of the OSI model.

TELNET

Telnet provides terminal emulation over a network. "Terminal" means text-based VT100-style terminal access. Telnet servers listen on TCP port 23. For more than twenty years, they were the standard way to access an interactive command shell over a network.

FTP

FTP, File Transfer Protocol, is used to transfer files to and from servers. Like Telnet, it has no confidentiality or integrity, and should not be used to transfer sensitive data over insecure channels.

TFTP

TFTP, Trivial File Transfer Protocol, runs on UDP port 69. It provides a simpler way to transfer files and is often used for saving router configurations or "bootstrapping" (downloading an operating system) via a network by diskless workstations.

TFTP has no authentication or directory structure: Files are read from and written to one directory, usually called /tftpboot. Nor does it have confidentiality or integrity. Like Telnet and FTP, TFTP is not recommended for transferring sensitive data over an insecure channel.

SSH

Secure Shell (SSH) was designed as a secure replacement for Telnet, FTP, and the UNIX "R" commands (rlogin, rshell, etc). It provides confidentiality, integrity, and secure authentication, among other features. SSH includes SFTP (SSH FTP) and SCP (Secure Copy) for transferring files. Also, it can be used to securely tunnel other protocols, such as HTTP. SSH servers listen on TCP port 22 by default.

SMTP, POP, AND IMAP

SMTP is the Simple Mail Transfer Protocol, used to transfer email between servers. SMTP servers listen on TCP port 25. **POPv3** (Post Office Protocol) and **IMAP** (Internet Message Access Protocol), which provide client-server email access, use TCP ports 110 and 143, respectively.

DNS

DNS, the Domain Name System, is a distributed global hierarchical database that translates names to IP addresses, and vice versa. It uses both TCP and UDP: small answers via UDP port 53; large answers (such as zone transfers), via TCP port 53.

HTTP AND HTTPS

HTTP is the Hypertext Transfer Protocol, used for transfer of unencrypted web-based data. **HTTPS** (Hypertext Transfer Protocol Secure) transfers encrypted web-based data via **SSL/TLS** (see the section SSL/TLS). HTTP uses TCP port 80; HTTPS, TCP port 443. **HTML** (Hypertext Markup Language) is used to display web content.

LAN technologies and protocols

Local Area Network concepts focus on layer 1 through layer 3 technologies such as physical and logical network topologies and Ethernet.

ETHERNET

Ethernet is a dominant local area networking technology that transmits network data via frames. It originally used a physical bus topology, but later added support for physical star. Ethernet handles Layer 1 issues such as physical media as well as Layer 2 issues such as frames. It is baseband (one channel), so it must deal with collisions, where two nodes attempt to transmit data simultaneously, and other similar issues.

WAN technologies and protocols

Wide Area Network technologies are often used by ISPs and other “long haul” service providers, whose networks span cities and countries. Many of us have hands-on experience configuring LAN technologies, such as connecting Cat5 network cabling; it is rarer to have hands-on experience building WANs.

T1s, T3s, E1s, AND E3s

There are a number of international circuit standards. The most prevalent are T Carriers (United States) and E Carriers (Europe).

Fast Facts

Here is a summary of common circuits:

- **T1:** a dedicated 1.544-megabit circuit that carries 24 64-bit DS0 (Digital Signal 0) channels.
- **T3:** 28 bundled T1s, forming a 44.736-megabit circuit.
- **E1:** a dedicated 2.048-megabit circuit that carries 30 channels.
- **E3:** 16 bundled E1s, forming a 34.368-megabit circuit.

Note that the terms *DS1* (Digital Signal 1) and *T1* are often used interchangeably. However, DS1 involves the flow of bits (via any medium, such as copper, fiber, wireless), whereas a T1 is a copper telephone circuit that carries a DS1. The terms *T3* and *DS3* (Digital Signal 3) are also used interchangeably, with the same T1/DS1 distinction just noted.

FRAME RELAY

Frame Relay is a packet-switched Layer 2 WAN protocol that provides no error recovery and focuses on speed. Higher-layer protocols carried by Frame Relay, such as TCP/IP, can provide reliability.

Frame Relay multiplexes multiple logical connections over a single physical connection to create Virtual Circuits. This shared bandwidth model is an alternative to dedicated circuits like T1s. A **PVC** (Permanent Virtual Circuit) is always connected, analogously to an actual dedicated circuit like a T1. An **SVC** (Switched Virtual Circuit) sets up each “call,” transfers data, and terminates the connection after an idle timeout. Frame Relay is addressed locally via Data Link Connection Identifiers (DLCIs, pronounced “delsees”).

X.25

X.25 is an older packet-switched WAN protocol that provided a cost-effective way to transmit data over long distances in the 1970s through early 1990s, when the most common alternative method was a direct call via analog modem. X.25’s popularity has faded as the Internet has become ubiquitous.

ATM

ATM is a WAN technology that uses fixed-length cells. ATM cells are 53 bytes long, with a 5-byte header and 48-byte data portion.

ATM allows reliable network throughput compared to Ethernet. The answer to “How many Ethernet frames can I send per second?” is “It depends.” Normal Ethernet frames can range in size from under 100 bytes to over 1500. In contrast, all ATM cells are 53 bytes.

The Switched Multimegabit Data Service (**SMDS**) is older and similar to ATM; it too uses 53-byte cells.

NETWORK DEVICES

We now look at network devices ranging from Layer 1 hubs through application-layer proxy firewalls that operate up to Layer 7. Many of these network devices, such as routers, use dedicated routing protocols.

Repeaters and hubs

Repeaters and hubs are Layer 1 devices. A repeater receives bits on one port and “repeats” them out the other port. It has no understanding of protocols but simply repeats bits. Repeaters are often used to extend the length of a network.

A hub is a repeater with more than two ports. It receives bits on one port and repeats them across all other ports.

Bridges

Bridges and switches are Layer 2 devices. A bridge has two ports and connects network segments together. Each segment typically has multiple nodes, and the bridge learns the MAC addresses of nodes on either side. Traffic sent from two

nodes on the same side of the bridge will not be forwarded across it. Traffic sent from a node on one side to a node on the other side will be forwarded. The bridge provides traffic isolation and makes forwarding decisions by learning the MAC addresses of connected nodes.

In Figure 7.1, traffic sent from Computer 1 to Computer 2 does not forward across the bridge. Traffic sent from Computer 1 to Computer 3 does.

A bridge has two collision domains. A network protocol analyzer (informally called a “sniffer”), shown on the right side of the network in Figure 7.1, can sniff traffic sent to or from Computers 3 and 4, but not traffic from Computers 1 and 2 (unless sent to Computer 3 or 4).

Switches

A switch is a bridge with more than two ports. It is best practice to connect only one device per switch port. Otherwise, everything that is true about a bridge is also true about a switch.

Figure 7.2 shows a network switch, which provides traffic isolation by associating the MAC address of each computer and server with its port. Traffic sent between Computer 1 and Server 1 remains isolated to the switch ports of those computers only: A network sniffer running on Server 3 does not see it.

Routers

Routers are Layer 3 devices that direct traffic from one LAN to another. IP-based routers make routing decisions based on the IP address of the source and of the destination.

Firewalls

Firewalls filter traffic between networks. TCP/IP packet filter and stateful firewalls make decisions based on layers 3 and 4 (IP addresses and ports); proxy

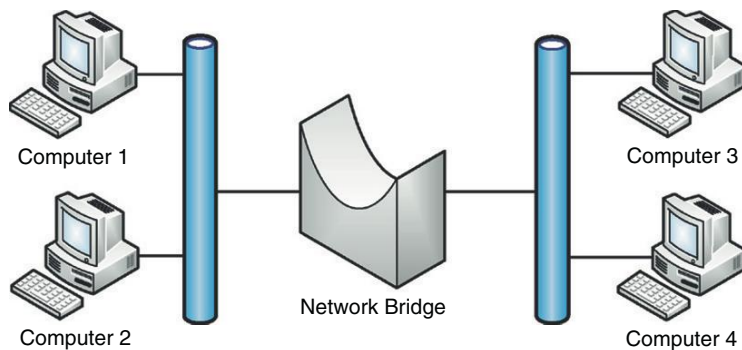


FIGURE 7.1
Network bridge.




FIGURE 7.2
Network switch.

firewalls can make decisions based on layers 5 through 7 as well. Generally, firewalls are multi-homed: They have multiple NICs connected to multiple different networks.

PACKET FILTER

A packet filter is a simple, fast firewall that has no concept of “state”: Each filtering decision must be made on the basis of a single packet. There is no way to refer to past packets in making current decisions.

STATEFUL FIREWALLS

Stateful firewalls have a state table that allows them to compare current and previous packets. They are slower than packet filters, but far more secure.

PROXY FIREWALLS

Proxies are firewalls that act as intermediary servers. Both packet filter and stateful firewalls pass traffic through or deny it: They are another hop along the route, passing on the TCP three-way handshake that takes place between the client and the server. Proxies terminate connections.

Application-Layer Proxy Firewalls

Application-layer proxies operate up to Layer 7. Unlike packet filter and stateful firewalls, which make decisions based on layers 3 and 4 only, they can make filtering decisions based on application-layer data, such as HTTP traffic as well.

DID YOU KNOW?

Application-layer proxies must understand the protocols that are proxied, so dedicated proxies are often required for each one: an FTP proxy for FTP traffic, an HTTP proxy for web traffic, and so on. This allows tighter control of filtering decisions. Instead of relying only on IP addresses and ports, an HTTP proxy can make decisions based on HTTP data, including web content. This allows sites to block access to explicit content.

Circuit-Level Proxies, Including SOCKS

Circuit-level proxies operate at Layer 5 (session), and are below application-layer proxies (at Layer 7). This allows them to filter more protocols: There is no need to understand each one; the application-layer data is simply passed along.

The most popular example of a circuit-level proxy, **SOCKS** has no need to understand application-layer protocols, so it can proxy many of them. Unlike its application-layer cousins, it cannot make fine-grained decisions: It does not understand application-layer protocols such as HTTP and thus cannot make filtering decisions based on application-layer data, such as explicit web content. SOCKS uses TCP port 1080.

Modem

A transmitting **Modem** (Modulator/Demodulator) modulates binary data into analog sound that can be carried on phone networks designed to transmit the human voice. A receiving modem demodulates the sound back into binary data. Modems are asynchronous: They do not operate with a clock signal.

SECURE COMMUNICATIONS

Protecting data in motion is a most complex challenge. The Internet provides cheap global communication...with little or no built-in confidentiality, integrity, or availability. We often must secure this data ourselves, and the concept of secure communications describes ways to accomplish that goal.

Authentication protocols and frameworks

An authentication protocol authenticates an identity claim over the network. Good security design assumes that a network eavesdropper may sniff all packets sent between a client and an authentication server: The protocol should remain secure. As we will see shortly, **PAP** fails this test, but **CHAP** and **EAP** pass.

PAP AND CHAP

The Password Authentication Protocol (PAP) is very weak. It sends the username and password in cleartext, which means that an attacker who is able to sniff the authentication process can launch a simple attack by replaying the username and password to log in. PAP is insecure and should not be used.

The Challenge-Handshake Authentication Protocol (CHAP) is more secure. It does not expose the cleartext password and so is not susceptible to replay attacks. CHAP relies on a shared secret, the password, which is securely created (e.g., during account enrollment) and stored on the CHAP server. Since both the user and the CHAP server share a secret, they can use that secret to securely communicate.

802.1X AND EAP

802.1X, “Port Based Network Access Control,” includes the Extensible Authentication Protocol (EAP), which is an authentication framework that describes many specific authentication protocols. EAP is designed to provide authentication at Layer 2 (it is “port based,” like ports on a switch) before a node receives an IP address. It is available for both wired and wireless networks, but is most commonly deployed on WLANs.

Exam Warning

Do not confuse 802.1X (EAP) with 802.11 (Wireless).

An EAP client, called a supplicant, requests authentication to a server, called an authenticator.

Fast Facts

There are many types of EAP; we will focus on **LEAP**, **EAP-TLS**, **EAP-TTLS**, and **PEAP**.

- LEAP (Lightweight Extensible Authentication Protocol) is a Cisco-proprietary protocol that was released before 802.1X was finalized. It has significant security flaws, and should not be used.
- EAP-TLS (EAP–Transport Layer Security) uses PKI, requiring both server-side and client-side certificates. It establishes a secure TLS tunnel for authentication. EAP-TLS is very secure because of PKI, but it is complex and costly for the same reason. The other major versions of EAP attempt to create the same TLS tunnel without requiring a client-side certificate.
- EAP-TTLS (EAP–Tunneled Transport Layer Security), developed by Funk Software and Certicom, simplifies EAP-TLS by dropping the client-side certificate requirement, allowing other methods (such as password) for authentication of clients. It is thus easier to deploy than EAP-TLS, but less secure.
- PEAP (Protected EAP) was jointly developed by Cisco Systems, Microsoft, and RSA Security. It is similar to EAP-TTLS (and may be considered a competitor) in that it does not require client-side certificates.

VPN

Virtual Private Networks (VPNs) secure data sent via insecure networks such as the Internet. The goal is to provide, virtually, the privacy provided by a circuit such as a T1. The nuts and bolts of VPNs involve authentication for security,

cryptographic hashes such as SHA-1 for integrity, and ciphers such as AES for confidentiality.

SLIP AND PPP

The Serial Line Internet Protocol (**SLIP**) resides at Layer 2 and provides IP connectivity via asynchronous connections such as serial lines and modems. When first introduced in 1988, it allowed routing packets via modem links for the first time (modems had until then been primarily used for nonrouted terminal access). SLIP is a bare-bones protocol that provides no built-in confidentiality, integrity, or authentication. It has largely faded from use, superseded by **PPP**.

As just mentioned, PPP (Point-to-Point Protocol), a Layer 2 resident, has largely replaced SLIP. It adds confidentiality, integrity, and authentication via point-to-point links, supporting synchronous links (e.g., T1) and asynchronous links (e.g., modems).

IPsec

IPv4 has no built-in confidentiality; it is the job of higher-layer protocols such as TLS to provide security. To address the lack of security at this layer, **IPsec** (Internet Protocol Security) was designed. It provides confidentiality, integrity, and authentication via encryption for IPv6. IPsec, which has been ported to IPv4, is a suite of protocols, the major two being Encapsulating Security Protocol (ESP) and Authentication Header (AH). Each has an IP protocol number: ESP, 50; AH, 51.

SSL AND TLS

Secure Sockets Layer (SSL) was designed to protect HTTP (HyperText Transfer Protocol) data: HTTPS uses TCP port 443. TLS (Transport Layer Security) is the latest SSL version, equivalent to version 3.1. The current TLS version is 1.2, described in RFC 5246 (see <http://tools.ietf.org/html/rfc5246>).

Though initially web-focused, SSL or TLS can encrypt many types of data and can tunnel other IP protocols to form VPN connections. SSL VPNs can be simpler than their IPsec equivalents: IPsec makes fundamental changes to IP networking, so its software installation changes the operating system (which requires super-user privileges). SSL client software does not require altering the operating system. Also, SSL is much simpler than IPsec to firewall.

Wireless Local Area Networks

Wireless Local Area Networks (WLANs) transmit information via electromagnetic waves (e.g., radio) or light. Historically, wireless data networks were very insecure, often relying on the (perceived) difficulty of attacking the confidentiality or integrity of the traffic. This perception was and is usually misplaced. The most common form of wireless data networking is 802.11, and the first 802.11 standard with reasonable security is 802.11i.

FHSS, DSSS, AND OFDM

Frequency Hopping Spread Spectrum (**FHSS**) and Direct Sequence Spread Spectrum (**DSSS**) are two methods for sending traffic via a radio band. Some bands, like the 2.4-GHz ISM, can become quite polluted with interference: Bluetooth, some cordless phones, some 802.11 wireless baby monitors, and even microwave ovens can interfere with this band. Both DSSS and FHSS are designed to maximize throughput while minimizing the effects of interference.

DSSS uses the entire band at once, “spreading” the signal throughout it. FHSS uses a number of small frequency channels throughout the band and “hops” through them in pseudorandom order.

Orthogonal Frequency-Division Multiplexing (**OFDM**) is a newer multiplexing method. It allows simultaneous transmission using multiple independent wireless frequencies that do not interfere with each other.

802.11

802.11 wireless has many standards, using various frequencies and speeds. The original mode, simply referred to as 802.11 (sometimes 802.11-1997, based on the year it was created), operated at 2 megabits per second (mbps) using the 2.4-GHz frequency; it was quickly supplanted by 802.11b, at 11 mbps. 802.11g was designed to be backward compatible with 802.11b devices, offering speeds up to 54 mbps using the 2.4-GHz frequency. 802.11a offers the same top speed, using the 5-GHz frequency.

802.11n uses both 2.4-GHz and 5-GHz frequencies, and takes advantage of multiple antennas with multiple-input multiple-output (MIMO). This allows speeds of 144 mbps and beyond. Table 7.3 summarizes the major types of 802.11 wireless.

The 2.4-GHz frequency can be quite crowded. The 5-GHz frequency is usually less so, and often does not have as much interference. It is a higher frequency with shorter waves, so it does not penetrate walls and other obstructions as easily as the longer 2.4-GHz waves can.

WEP

WEP, the Wired Equivalent Privacy protocol, was an early attempt (first ratified in 1999) to provide 802.11 wireless security. It has proven to be critically weak:

Table 7.3 Types of 802.11 Wireless

Type	Top Speed	Frequency
802.11	2 mbps	2.4 GHz
802.11a	54 mbps	5 GHz
802.11b	11 mbps	2.4 GHz
802.11g	54 mbps	2.4 GHz
802.11n	144 mbps+	2.4 GHz and/or 5 GHz

New attacks can break any WEP key in minutes. For this reason, WEP provides little integrity or confidentiality protection. Indeed, it is considered broken and its use is strongly discouraged. 802.11i and/or other encryption methods such as VPN should be used instead.

802.11i

802.11i, the first 802.11 standard to provide reasonable security, describes a Robust Security Network (**RSN**), which allows pluggable authentication modules. RSN allows changes to cryptographic ciphers as new vulnerabilities are discovered.

Crunch Time

RSN is also known as **WPA2** (Wi-Fi Protected Access 2), which is a full implementation of 802.11i. By default WPA2 uses AES encryption to provide confidentiality, and **CCMP** (Counter Mode CBC MAC Protocol) to create a Message Integrity Check (**MIC**) for integrity. It may (optionally) use the less secure **RC4** (Rivest Cipher 4) and **TKIP** (Temporal Key Integrity Protocol) ciphers to provide confidentiality and integrity, respectively.

The less secure **WPA** (without the “2”) was designed for access points that lack the power to implement the full 802.11i standard, so providing a better security alternative to WEP. It uses RC4 for confidentiality and TKIP for integrity. WPA2 is recommended over WPA.

Remote access

In this age of telecommuting and the mobile workforce, secure remote access is a critical control. This includes connecting mobile users via methods such as DSL and cable modem, security mechanisms such as callback, and newer concerns such as instant messaging and remote meeting technology.

ISDN

Integrated Services Digital Network (**ISDN**) was an earlier attempt to provide digital service via “copper pair,” the Plain Old Telephone Service (POTS) prevalent in homes and small offices around the world. The telephone cables, which connect telecom central offices to customer, are called the “last mile”; providing high-speed digital service via the (historically copper pair) last mile has been a longstanding challenge.

DSL

Digital Subscriber Line (**DSL**) is a “last mile” solution similar to ISDN: It uses existing copper pairs to provide homes and small offices with digital service. Compared to ISDN, DSL is more widely used because of its higher speeds, which can reach 10 megabits and more.

Table 7.4 DSL Speed and Distances

Type	Download Speed	Upload Speed	Distance from CO
ADSL	1.5–9 mbps	16–640 Kbps	18,000 feet
SDSL	1.544 mbps	1.544 mbps	10,000 feet
HDSL	1.544 mbps	1.544 mbps	10,000 feet
VDSL	20–50+ mbps	Up to 20 mbps	<5,000 feet

Source: Cisco Press, *DSL and Cable Modem Networks*. URL: www.ciscopress.com/articles/article.asp?p=31289 (accessed July 23, 2010).

Common DSL types are Symmetric Digital Subscriber Line (SDSL; also called Single-Line DSL), with matching upload and download speeds; Asymmetric Digital Subscriber Line (ADSL), featuring faster downloads than uploads; and Very High Rate Digital Subscriber Line (VDSL), featuring much greater asymmetric speeds. Another option is HDSL (High Data Rate DSL), which matches SDSL speeds using two pairs of copper; it is used to provide inexpensive T1 service.

One advantage of ADSL is that it allows simultaneous use of a POTS line, often filtered from the DSL traffic. As a general rule, the closer a site is to the Central Office (CO), the faster the available service.

Table 7.4 summarizes the speeds and modes of DSL.

CABLE MODEMS

Cable modems are used by Cable TV companies to provide Internet access via broadband cable TV. Cable TV access is not ubiquitous, but is available in most large towns and cities in industrialized areas. Broadband, unlike baseband, has multiple channels (similar to TV channels), so dedicating bandwidth for network services requires dedicating channels for that purpose. Cable modems are a compelling “last mile” solution: The providers have already invested in connecting the last mile, and Internet service offers another revenue stream based on that investment.

Unlike DSL, cable modem bandwidth is typically shared with neighbors on the same network segment.

SUMMARY OF EXAM OBJECTIVES

The domain of Telecommunications and Network Security is a large and complex one, requiring broad and sometimes deep understanding of thorny technical issues. Because our modern world relies on networks, they must be kept secure. It is important to understand not only why we use concepts like packet-switched networks and the OSI model but also how we implement them.

TOP FIVE TOUGHEST QUESTIONS

1. What is the difference between a stateful firewall and a proxy firewall?
 - A. Stateful firewalls have state; proxy firewalls do not
 - B. Proxy firewalls have state; stateful firewalls do not
 - C. Stateful firewalls terminate connections; proxy firewalls do not
 - D. Proxy firewalls terminate connections; stateful firewalls do not
2. Which is the most secure type of EAP?
 - A. EAP-TLS
 - B. EAP-TTLS
 - C. LEAP
 - D. PEAP
3. Which WAN protocol has no error recovery, relying on higher-level protocols to provide reliability?
 - A. ATM
 - B. Frame Relay
 - C. SMDS
 - D. X.25
4. Which is the most secure type of firewall?
 - A. Packet filter
 - B. Stateful
 - C. Circuit-level proxy
 - D. Application-layer proxy
5. Which term describes accessing an IPv6 network via an IPv4 network?
 - A. CIDR
 - B. NAT
 - C. Translation
 - D. Tunneling

Answers

1. Correct Answer and Explanation: **D**. Answer **D** is correct; proxy firewalls terminate connections, including the three-way TCP handshake. Stateful firewalls do not.
Incorrect Answers and Explanations: **A**, **B**, and **C**. Answers **A**, **B**, and **C** are incorrect. Stateful firewalls have a state table, so **A** is true, but a weaker answer than **D**: The key difference is that proxies terminate connections. **B** is the reverse of **A** and therefore false. **C** is the reverse of the correct answer **D** and therefore false.
2. Correct Answer and Explanation: **A**. Answer **A** is correct; EAP-TLS is the most secure (and costly) EAP type because it requires both server-side and client-side certificates.
Incorrect Answers and Explanations: **B**, **C**, and **D**. Answers **B**, **C**, and **D** are incorrect. EAP-TTLS and PEAP are similar and do not require client-side certificates. LEAP is a Cisco-proprietary protocol that does not require client-side certificates and has fundamental security weaknesses.

- 3.** Correct Answer and Explanation: **B.** Answer **B** is correct; Frame Relay is a packet-switched Layer 2 WAN protocol that features no error recovery.
Incorrect Answers and Explanations: **A, C, and D.** Answers **A, C, and D** are incorrect. ATM and SMDS are cell-based WAN protocols that provide error correction. X.25 is a packet-switched protocol that is similar to Frame Relay but features error recovery.
- 4.** Correct Answer and Explanation: **D.** Answer **D** is correct; application-layer firewalls are the most secure: They have the ability to filter based on OSI Layers 3 through 7.
Incorrect Answers and Explanations: **A, B, and C.** Answers **A, B, and C** are incorrect. All are firewalls. A packet filter is the least secure of the four because of the lack of state. A stateful firewall is more secure than a packet filter, but its decisions are limited to Layers 3 and 4. Circuit-level proxy firewalls operate at Layer 5, and cannot filter based on application-layer data.
- 5.** Correct Answer and Explanation: **D.** Answer **D** is correct; accessing an IPv6 network via an IPv4 network is called tunneling.
Incorrect Answers and Explanations: **A, B, and C.** Answers **A, B, and C** are incorrect. CIDR (Classless Inter-Domain Routing) is a way to create flexible subnets. NAT (Network Address Translation) translates one IP address into another. **C, Translation,** is a distracter answer.

This page intentionally left blank

CHAPTER 8

Domain 8: Application Development Security

129

Exam Objectives in this Chapter

- Programming concepts
- Application development methods
- Object-oriented programming
- Software vulnerabilities, testing, and assurance
- Databases

INTRODUCTION

Software is everywhere: not only in our computers but in our houses, our cars, and our medical devices. And all software programmers make mistakes. As our software has grown in complexity, the number of mistakes has grown apace.

Developing software that is robust and secure is critical: This chapter will show how to do that. We will cover programming fundamentals such as compiled versus interpreted languages as well as procedural and object-oriented programming languages. We will discuss application development models such as the Waterfall Model, the Spiral Model, eXtreme Programming (XP), and others. We will describe common software vulnerabilities, ways to test for them, and maturity frameworks to assess the maturity of the programming process and provide ways to improve it.

PROGRAMMING CONCEPTS

We begin by understanding some cornerstone programming concepts. As computers have become more powerful and ubiquitous, the process and methods used to create software have grown and changed.

Machine code, source code, and assemblers

Machine code (also called machine language) is software that is executed directly by the CPU. Machine code is CPU-dependent; it is a series of ones and zeroes that translate to instructions that the CPU understands. **Source code** is

computer programming language instructions written in text that must be translated into machine code before execution by the CPU. High-level languages contain English-like instructions such as `printf` (print formatted).

Assembly language is a low-level computer programming language. Its instructions are short mnemonics, such as `ADD`, `SUB` (subtract), and `JMP` (jump), that match machine language instructions. An assembler converts assembly language into machine language. A disassembler converts machine language into assembly.

Compilers, interpreters, and bytecode

Compilers take source code, such as C or Basic, and compile it into machine code, typically saved in executable form. **Interpreted** languages differ from compiled languages: Interpreted code (e.g., shell code) is compiled on the fly each time the program is run. If an interpreted program is run 100 times, it will be compiled 100 times (whereas a compiled program is compiled only once).

Publicly released software

Once programmed, publicly released software may come in different forms (e.g., with or without the accompanying source code) and released under a variety of licenses.

OPEN- AND CLOSED-SOURCE SOFTWARE

Closed-source software is typically released in executable form: The source code is kept confidential. Examples include Oracle and Microsoft Windows 7. Open-source software publishes source code publicly, allowing anyone to inspect, modify, or compile it. Examples include Ubuntu Linux and the Apache web server. Proprietary software is subject to intellectual property protections such as patents or copyrights. The terms “closed-source” and “proprietary” are sometimes used synonymously, but that is not always correct: Some open-source software is also proprietary.

FREE SOFTWARE, SHAREWARE, AND CRIPPLEWARE

Free software is a controversial term defined differently by different groups. “Free” may mean free of charge to use (sometimes called “free as in beer”), or it may mean that the user is free to use the software in any way he or she chooses, including modifying it (sometimes called “free as in liberty”). The two types are called *gratis* and *libre*, respectively. The confusion derives from the fact that “free” carries multiple meanings in English. Software that is both *gratis* and *libre* is sometimes called *free*² (free squared).

Freeware is “free as in beer” (*gratis*) that is, free of charge. **Shareware** is fully functional proprietary software that may be initially used free of charge. If the user continues to use it for a period of time specified by the license (such as 30 days), the Shareware license typically requires payment. **Crippleware** is partially functioning proprietary software, often with key features disabled. The user is typically required to make a payment to unlock the full functionality.

APPLICATION DEVELOPMENT METHODS

As software has grown in complexity, programming has increasingly become a team effort. Team-based projects require project management: to provide a framework with deliverables and milestones, to divvy up tasks, direct team communication, evaluate and report progress, and (hopefully) deliver a final product.

Ultimately, large application development projects may closely resemble projects that have nothing to do with software, such as widget production or bridge building. Development methods such as the Waterfall and Spiral Models are often close cousins to nonprogramming models. They can be thought of as project management methods, with additional features to support code writing.

Waterfall model

The Waterfall Model is a linear application development model that uses rigid phases: When one phase ends, the next begins. Steps occur in sequence, and, if unmodified, the model does not allow developers to go back to previous steps (hence “waterfall”: Once water falls down, it cannot go back up).


Exam Warning

The phases in the Waterfall Model are not specifically testable: Learn the overall flow. Also, the model omits a critical final step: destruction. No development process that leads to an operational system with sensitive production data is truly complete until that system has been retired, the data archived, and the remaining data on the system securely destroyed.

The Modified Waterfall Model allows a return to a previous phase for verification or validation, ideally confined to connecting steps. Barry Boehm’s paper “A Spiral Model of Software Development and Enhancement” (see the next section) discusses a modified waterfall based on Royce’s paper, shown in Figure 8.1.

Spiral

The Spiral Model is designed to control risk. It repeats the steps of a project, starting with modest goals and expanding outward in ever wider spirals (called rounds). Each round constitutes a project and may follow a traditional software development methodology such as the Modified Waterfall. A risk analysis is performed at each round. Fundamental flaws in the project or process are more likely to be discovered in the earlier phases, resulting in simpler fixes. This lowers the overall risk of the project: Large risks should be identified and mitigated.

**FIGURE 8.1**

Modified Waterfall (Spiral) Development Model. Source: Boehm, Barry. *A Spiral Model of Software Development and Enhancement*. URL: <http://portal.acm.org/citation.cfm?id=12948> (accessed July 23, 2010).

eXtreme Programming

eXtreme Programming (XP) is an Agile Software development method that uses pairs of programmers working off a detailed specification. There is a high level of customer involvement.

eXtreme Programming improves a software project in five essential ways; communication, simplicity, feedback, respect, and courage. eXtreme programmers constantly communicate with their customers and fellow programmers. They keep their design simple and clean. They get feedback by testing their software starting on day one. They deliver the system to the customers as early as possible and implement changes as suggested.¹

(See www.extremeprogramming.org/rules.html for more information.)

Rapid Application Development

In Rapid Application Development (RAD) software is developed via the use of prototypes, “dummy” GUIs, back-end databases, and more. The goal is quickly meeting the system’s business need; technical concerns are secondary. The customer is heavily involved in the process.

SDLC

The Systems Development Life Cycle (SDLC; also called *Software Development Life Cycle* or simply *System Life Cycle*) is a system development model used throughout the IT industry. However, SDLC focuses on security when used in the context of the exam. Think of “our” SDLC as the “*Secure Systems Development Life Cycle*”: The security is implied.

Fast Facts

The following overview is summarized from NIST Special Publication 800-14²:

Prepare a Security Plan. Ensure that security is considered during all phases of the IT system life cycle, and that security activities are accomplished during each phase.

Initiation. Need for a system is expressed and its purpose is documented.

- *Conduct a Sensitivity Assessment:* Look at the security sensitivity of the system and the information to be processed.

Development/Acquisition. The system is designed, purchased, programmed, or developed

- *Determine Security Requirements:* Determine technical features (e.g., access controls), assurances (e.g., background checks for system developers), or operational practices (e.g., awareness and training).
- *Incorporate Security Requirements into Specifications:* Ensure that the previously gathered information is incorporated in the project plan
- *Obtain the System and Related Security Activities:* May include developing the system’s security features, monitoring the development process itself for security problems, responding to changes, and monitoring threats.

Continued

Implementation. Test and install the system.

- *Install/Turn on Controls:* A system often comes with security features disabled. These need to be enabled and configured.
- *Security Testing:* Used to certify a system; may include testing security management, physical facilities, personnel, procedures, the use of commercial or in-house services (such as networking services), and contingency planning.
- *Accreditation:* The formal authorization by the accrediting (management) official for system operation and an explicit acceptance of risk.

Operation/Maintenance. The system is modified by the addition of hardware and software and by other events.

- *Security Operations and Administration:* Examples include backups, training, managing cryptographic keys, user administration, and patching.
- *Operational Assurance:* Examines whether a system is operated according to its current security requirements.
- *Audits and Monitoring:* A system audit is a one-time or periodic event to evaluate security. Monitoring refers to an ongoing activity that examines either the system or the users.

Disposal. The secure decommissioning of a system.

- *Information:* Information may be moved to another system, archived, discarded, or destroyed.
- *Media Sanitization:* There are three general methods of purging media: overwriting, degaussing (for magnetic media only), and destruction.

OBJECT-ORIENTED PROGRAMMING

Object-Oriented Programming (OOP) uses an object metaphor to design and write computer programs. An object is a “black box” that can perform functions and that sends and receives messages. It contains data and methods (the functions they perform), and it provides encapsulation (also called data hiding): We do not know, from the outside, how the object performs its function. This provides security benefits: Users should not be exposed to unnecessary details. Examples of OOP languages include Java, C++, Smalltalk, and Ruby.

Cornerstone Object-Oriented Programming concepts

Cornerstone OOP concepts include objects, methods, messages, inheritance, delegation, polymorphism, and polyinstantiation. We will use an example object called Addy to illustrate them. Addy is an object that adds two integers; it is extremely simple, but has enough complexity to explain core OOP concepts. Addy inherits an understanding of numbers and math from his parent class (called mathematical operators). One specific object is called an instance. Note that objects may inherit from other objects in addition to classes.

FIGURE 8.2
The Addy object.



In our case, the programmer simply needs to program Addy to support the method of addition (inheritance takes care of everything else Addy must know). Figure 8.2 shows Addy adding two numbers.

1 + 2 is the input message; 3 is the output message. Addy also supports delegation: If he doesn't know how to perform a requested function, he can delegate that request to another object (called Subby in Figure 8.3).




FIGURE 8.3
Delegation.

Addy also supports polymorphism (based on the Greek roots *poly* and *morph*, meaning *many* and *forms*, respectively): He has the ability to overload his plus (+) operator, performing different methods depending on the context of the input message. For example: Addy adds when the input message contains number + number; polymorphism allows him to concatenate two strings when the input message contains string + string, as shown in Figure 8.4.

Finally, *polyinstantiation* means “many instances,” or two instances (specific objects) with the same name that contain different data. Figure 8.5 shows polyinstantiated Addy objects: two objects with the same name but different data.



FIGURE 8.4
Polymorphism.

Note that these are two separate objects. Also, to a secret-cleared subject, the Addy object with secret data is the only Addy object known.




FIGURE 8.5
Polyinstantiation.

Fast Facts

Here is a summary of Object Oriented Programming concepts as illustrated by Addy:

- *Object*: Addy.
- *Class*: mathematical operators.
- *Method*: addition.
- *Inheritance*: Addy inherits an understanding of numbers and math from his parent class mathematical operators. The programmer simply needs to program Addy to support the method of addition.
- *Example input message*: 1 + 2.
- *Example output message*: 3.
- *Polymorphism*: Addy can change behavior based on the context of the input, overloading the + to perform addition or concatenation depending on the context.
- *Polyinstantiation*: Two Addy objects (secret and top secret), with different data.

Object Request Brokers

As we have seen, mature objects are designed to be reused: They lower risk and development costs. Object Request Brokers (**ORBs**) can be used to locate objects: They act as object search engines. ORBs are middleware, connecting programs to programs. Common object brokers include COM, DCOM, and CORBA.

COM AND DCOM

Two object broker technologies by Microsoft are Component Object Model (**COM**) and Distributed Component Object Model (**DCOM**). COM locates objects on a local system; DCOM can also locate objects over a network.

COM allows objects written in different OOP languages to communicate, where objects written in C++ send messages to objects written in Java, for example. COM is designed to hide the details of any individual object, instead focusing on the object's capabilities.

DCOM is a networked sequel to COM:

Microsoft® Distributed COM (DCOM) extends the Component Object Model (COM) to support communication among objects on different computers—on a LAN, a WAN, or even the Internet. With DCOM, your application can be distributed at locations that make the most sense to your customer and to the application.³

DCOM includes Object Linking and Embedding (**OLE**), which is a way to link documents together. Both COM and DCOM are being supplanted by Microsoft .NET, which can interoperate with DCOM, but offers advanced functionality to both DCOM and COM.

CORBA

CORBA (Common Object Request Broker Architecture) is an open vendor-neutral networked object broker framework developed by the Object Management Group (OMG). CORBA competes with Microsoft's proprietary DCOM. Its objects communicate via a message interface, described by the following quote from the Interface Definition Language (**IDL**). (See www.corba.org for more information.)

The essence of CORBA, beyond its being a networked object broker, is the separation of the interface (the syntax for communicating with an object) from the instance (the specific object): The interface to each object is defined very strictly. In contrast, the implementation of an object—its running code, and its data—is hidden from the rest of the system (that is, encapsulated) behind a boundary that the client may not cross. Clients access objects only through their advertised interface, invoking only those operations that the object exposes through its IDL interface, with only those parameters (input and output) that are included in the invocation.⁴

SOFTWARE VULNERABILITIES, TESTING, AND ASSURANCE

Once the project is under way and the software has been programmed, the next step is testing the software, focusing on the confidentiality, integrity, and availability of the system, the application, and the data processed by the application. Special care must be given to the discovery of software vulnerabilities that could lead to data or system compromise. Finally, organizations need to be able to gauge the effectiveness of their software creation process and identify ways to improve it.

Software vulnerabilities

That programmers make mistakes has been true since the advent of computer programming. In *Code Complete*, Steve McConnell says “experience suggests that there are 15 to 50 errors per 1000 lines of delivered code.”⁵ A thousand lines of code is sometimes called a KLOC (*K* stands for thousand). This number can be lowered by following a formal application maturity framework model. Watts S. Humphrey, a fellow at Carnegie Mellon University’s Software Engineering Institute, claims that organizations that follow the SEI Capability Maturity Model (CMM; see the section on the Software Capability Maturity Model) can lower the number of errors to one in every KLOC.⁶

TYPES OF SOFTWARE VULNERABILITIES

This section will briefly describe common application vulnerabilities. More technical details on vulnerabilities such as buffer overflows were discussed in Chapter 5. An additional source of up-to-date vulnerabilities can be found in “2010 CWE/SANS Top 25 Most Dangerous Programming Errors,” available at <http://cwe.mitre.org/top25/>. The following summary is based on this list. CWE (Common Weakness Enumeration) is a dictionary of software vulnerabilities by MITRE (see <http://cwe.mitre.org/>). SANS is the SANS Institute (see www.sans.org).⁷

- *Hard-Coded Credentials*: Backdoor username/passwords left by programmers in production code.
- *Buffer Overflow*: Occurs when a programmer does not perform variable bounds checking.
- *SQL Injection*: Manipulation of a back-end SQL server via a front-end web server.
- *Directory Path Traversal*: Escaping from the root of a web server (such as `/var/www`) into the regular file system by referencing directories such as `../`.
- *PHP Remote File Inclusion (RFI)*: Altering normal PHP URLs and variables, such as `http://good.example.com?file=readme.txt`, to include and execute remote content (e.g., `http://good.example.com?file=http://evil.example.com/bad.php`).
- *Cross-Site Scripting (XSS)*: Third-party execution of web scripting languages such as Javascript within the security context of a trusted site.
- *Cross-Site Request Forgery (CSRF, or sometimes XSRF)*: Third-party redirect of static content within the security context of a trusted site.

Cross-Site Scripting (XSS) and Cross-Site Request Forgery (CSRF) are often confused. They are both web attacks; the difference is that XSS executes a script in a trusted context:

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

This code would pop up a harmless XSS Test! alert. A real attack would include more Javascript, often stealing cookies or authentication credentials.

CSRF can trick a user into processing a URL (sometimes by embedding the URL in an HTML image tag) that performs a malicious act—for example, tricking a white hat into rendering the following image tag:

```

```

Disclosure

Disclosure describes the actions taken by a security researcher after discovering a software vulnerability. This topic has proven controversial: What actions should you take if, in well-known software such as the Apache web server or Microsoft's IIS (Internet Information Services) web server, you discover a flaw?

Full disclosure is the practice of releasing vulnerability details publicly. The rationale is this: If the bad guys may already have the information, everyone should have it. This ensures that the white hats know about the vulnerability and will pressure the vendor to patch it. Advocates argue that vulnerable software should be fixed as quickly as possible; relying on a (perceived) lack of awareness of the vulnerability amounts to "security through obscurity," which many argue is ineffective. The Full Disclosure mailing list (see <http://seclists.org/fulldisclosure/>) is dedicated to full disclosure.

Responsible disclosure is the private sharing of vulnerability information with a vendor, withholding public release until a patch is available. This is generally considered to be the ethical option. Other options fall between full and responsible disclosure, including privately sharing vulnerability information with a vendor but including a deadline, such as "I will post the vulnerability details publicly in three months, or after you release a patch, whichever comes first."

Software Capability Maturity Model

The Software Capability Maturity Model (CMM) is a framework for evaluating and improving the software development process. It was developed by the Software Engineering Institute (SEI) of Carnegie Mellon University (CMU).

The goal of CMM is to develop a methodical framework for creating quality software that allows measurable and repeatable results:

Even in undisciplined organizations, however, some individual software projects produce excellent results. When such projects succeed, it is generally through the heroic efforts of a dedicated team, rather than

through repeating the proven methods of an organization with a mature software process. In the absence of an organization-wide software process, repeating results depends entirely on having the same individuals available for the next project. Success that rests solely on the availability of specific individuals provides no basis for long-term productivity and quality improvement throughout an organization. Continuous improvement can occur only through focused and sustained effort towards building a process infrastructure of effective software engineering and management practices.⁸

Fast Facts

The five levels of CMM are described as quoted here (see www.sei.cmu.edu/reports/93tr024.pdf):

1. *Initial*: The software process is characterized as ad hoc, and occasionally even chaotic. Few processes are defined, and success depends on individual effort.
2. *Repeatable*: Basic project management processes are established to track cost, schedule, and functionality. The necessary process discipline is in place to repeat earlier successes on projects with similar applications.
3. *Defined*: The software process for both management and engineering activities is documented, standardized, and integrated into a standard software process for the organization. Projects use an approved, tailored version of the organization's standard software process for developing and maintaining software.
4. *Managed*: Detailed measures of the software process and product quality are collected, analyzed, and used to control the process. Both the software process and products are quantitatively understood and controlled.
5. *Optimizing*: Continual process improvement is enabled by quantitative feedback from the process and from piloting innovative ideas and technologies.⁸

DATABASES

A database is a structured collection of related data that allows queries (searches), insertions (updates), deletions, and many other functions. It is managed by the Database Management System (DBMS), which controls all access to it and enforces database security. Databases are overseen by Database Administrators (DBAs). They may be searched with a database query language, such as the Structured Query Language (SQL). Typical database security issues include confidentiality and integrity of the stored data. Integrity is a primary concern when replicated databases are updated.

Relational databases

The most common modern database is relational, which means that it contains two-dimensional tables of related data (hence *relational*). A database table is also called a relation. Tables have rows and columns: A row is a database record, called a **tuple**; a column is called an **attribute**. A single cell (the intersection of a

Table 8.1 Relational Database Employee Table

SSN	Name	Title
133-73-1337	J.F. Sebastian	Designer
343-53-4334	Eldon Tyrell	Doctor
425-22-8422	Gaff	Detective
737-54-2268	Rick Deckard	Detective
990-69-4771	Hannibal Chew	Engineer

row and a column) in a database is called a **value**. Relational databases require a unique value called the primary key in each tuple in a table. Table 8.1 shows a relational database employee table, sorted by the primary key (Social Security Number, SSN).

Table 8.1 attributes are SSN, Name, and Title. Tuples include each row: 133-73-1337, 343-53-4334, and so forth. “Gaff” is an example of a value (cell). Candidate keys are any attribute (column) in the table with unique values: In Table 8.1 these include SSN and Name; SSN was selected as the primary key because it is truly unique (two employees can have the same name, but not the same SSN). Two tables in a relational database may be joined by the primary key.

FOREIGN KEYS

A foreign key in a relational database table matches a primary key in the parent database. Note that the foreign key is the local table’s primary key: It is called foreign when referring to the parent. Table 8.2 is an HR database table listing employees’ vacation time (in days) and sick time (in days); its foreign is the SSN. This database table may be joined to the parent (employee) database table by connecting its foreign key to the employee table’s primary key.

REFERENTIAL, SEMANTIC, AND ENTITY INTEGRITY

Databases must ensure the integrity of their table data: This is called data integrity, discussed in the section Database Integrity to come. There are three additional specific integrity issues that must be addressed beyond the correctness of the data itself: Referential, Semantic, and Entity. These are closely tied to the logical operations of the DBMS.

Table 8.2 HR Database Table

SSN	Vacation Time	Sick Time
133-73-1337	15 days	20 days
343-53-4334	60 days	90 days
425-22-8422	10 days	15 days
737-54-2268	3 days	1 day
990-69-4771	15 days	5 days

Crunch Time

Referential integrity means that every foreign key in a secondary table matches a primary key that is in the parent table: If this is not true, referential integrity has been broken. Semantic integrity means that every attribute

(column) value is consistent with the attribute data type. Entity integrity means that every tuple has a unique primary key that is not null.

Table 8.3 Database Table Lacking Integrity

SSN	Vacation Time	Sick Time
467-51-9732	7 days	14 days
737-54-2268	3 days	Nexus 6
133-73-1337	16 days	22 days
133-73-1337	15 days	20 days

The HR database table (Table 8.2) has referential, semantic, and entity integrity. Table 8.3, on the other hand, has multiple problems: One tuple violates referential integrity, one violates semantic integrity, and two violate entity integrity.

The tuple with the foreign key 467-51-9732 has no matching entry in the employee database table. This breaks referential integrity, as there is no way to link the entry to a name or title. Cell “Nexus 6” violates semantic integrity because the sick time attribute requires values of days, and Nexus 6 is not such a value. Finally, the last two tuples have the same primary key (primary to this table; foreign to the parent employees table), which breaks entity integrity.

DATABASE NORMALIZATION

Database normalization seeks to make the data in a database table logically concise, organized, and consistent. It removes redundant data and improves database integrity and availability.

DID YOU KNOW?

Normalization has three rules, called forms (see www.informit.com/articles/article.aspx?p=30646 for more information)¹⁰:

1. First Normal Form (1NF): Divide data into tables.
2. Second Normal Form (2NF): Move data that is partially dependent on the primary key to another table. The HR Database (Table 8.2) is an example of 2NF.
3. Third Normal Form (3NF): Remove data that is not dependent on the primary key.⁹

Table 8.4 Employee Table Database View: “Detective”

SSN	Name	Title
425-22-8422	Gaff	Detective
737-54-2268	Rick Deckard	Detective

DATABASE VIEWS

Database tables may be queried; the results of a query are called a database view. Views may be used to provide a **constrained user interface**: For example, nonmanagement employees can be shown their individual records only via database views. Table 8.4 shows the database view resulting from querying the employee table Title attribute with a string of Detective. While employees of the HR department may be able to view the entire employee table, this view may be authorized only for the captain of detectives, for example.

DATABASE QUERY LANGUAGES

Database query languages allow the creation of database tables, read/write access to those tables, and many other functions. They have at least two subsets of commands: Data Definition Language (**DDL**) and Data Manipulation Language (**DML**). DDL is used to create, modify, and delete tables. DML is used to query and update table data.

Database integrity

In addition to the previously discussed relational database integrity issues of semantic, referential, and entity integrity, databases must also ensure data integrity—that is, the integrity of the table entries. This treats integrity as a more general issue: mitigating unauthorized data modifications. The primary challenge is simultaneous attempted modifications of data. A database server typically runs multiple threads (lightweight processes), each capable of altering data. What happens if two threads attempt to alter the same record?

DBMSs may attempt to commit updates—that is, make pending changes permanent. If the update commit is unsuccessful, the DBMS can roll back (or abort) and restore from a savepoint (i.e., a clean snapshot of the table).

A database journal is a log of all database transactions. Should a database become corrupted, it can be reverted to a backup copy; subsequent transactions can then be “replayed” from the journal, restoring database integrity.

Database Replication and Shadowing

Databases may be highly available (HA)—in other words, replicated with multiple servers containing multiple table copies. Integrity is the primary concern with replicated databases: If a record is updated in one table, it must be simultaneously

updated in all tables. What happens if two processes attempt to update the same tuple simultaneously on two different servers? Both cannot be successful; this would violate the tuple's integrity.

Database replication mirrors a live database, allowing simultaneous client-generated reads and writes to multiple replicated databases. Replicated databases pose additional integrity challenges. A two-phase (or multi-phase) commit can be used to ensure integrity before committing in the following way. The DBMS requests a vote. If the DBMS on each server agrees to commit, the changes are made permanent. If any DBMSs disagree, the vote fails and the changes are not committed (made permanent).

A shadow database is similar to a replicated database, with one key difference: It mirrors all changes made to a primary database but clients do not access it. Unlike replicated databases, the shadow database is one-way (data flows from primary to shadow): It serves as a live data backup of the primary.

SUMMARY OF EXAM OBJECTIVES

We live in an increasingly computerized world, and software is everywhere. Data confidentiality, integrity, and availability are critical, as is the normal functionality (availability) of the software that processes it. This domain has shown how software works, and has highlighted the challenges programmers face in writing error-free code that can protect both the data and itself in the face of attacks.

The use of a formal methodology for developing software followed by a rigorous testing regimen is best practice. We have seen that a software development maturity model such as the Capability Maturity Model (CMM) can dramatically lower the number of errors programmers make. The five steps of the CMM mimic the process most programming organizations follow, from informal to mature, always seeking improvement: initial, repeatable, defined, managed, and optimized.

TOP FIVE TOUGHEST QUESTIONS

1. Which type of database language is used to create, modify, and/or delete tables?
 - A. Data Definition Language (DDL)
 - B. Data Manipulation Language (DML)
 - C. Database Management System (DBMS)
 - D. Structured Query Language (SQL)
2. A database contains an entry with an empty primary key. What database concept has been violated?
 - A. Entity Integrity
 - B. Normalization
 - C. Referential Integrity
 - D. Semantic Integrity

3. Which vulnerability allows a third party to redirect static content within the security context of a trusted site?
 - A. Cross-Site Request Forgery (CSRF)
 - B. Cross-Site Scripting (XSS)
 - C. PHP Remote File Inclusion (RFI)
 - D. SQL Injection
4. Which language allows CORBA (Common Object Request Broker Architecture) objects to communicate via a message interface?
 - A. Distributed Component Object Model (DCOM)
 - B. Interface Definition Language (IDL)
 - C. Object Linking and Embedding (OLE)
 - D. Object Management Guidelines (OMG)
5. Which database high-availability option allows multiple clients to access multiple database servers simultaneously?
 - A. Database commit
 - B. Database journal
 - C. Replicated database
 - D. Shadow database

Answers

1. Correct Answer and Explanation: A. Answer A is correct; Data Definition Language (DDL) is used to create, modify, and delete tables.
Incorrect Answers and Explanations: B, C, and D. Answers B, C, and D are incorrect. Data Manipulation Language (DML) is used to create, modify, and delete tables. Data Manipulation Language (DML) is used to query and update data stored in the tables. Database Management System (DBMS) manages the database system and provides security features. Structured Query Language (SQL) is a database query language that includes both DDL and DML. DDL is more specific than SQL, so it is a better answer for this question.
2. Correct Answer and Explanation: A. Answer A is correct; entity integrity means that each tuple has a unique primary key that is not null.
Incorrect Answers and Explanations: B, C, and D. Answers B, C, and D are incorrect. Normalization makes the data in a database table logically concise, organized, and consistent. Referential integrity means that every foreign key in a secondary table matches a primary key in the parent table: If this is not true, referential integrity has been broken. Semantic integrity means that each attribute (column) value is consistent with the attribute data type.
3. Correct Answer and Explanation: A. Answer A is correct; Cross-Site Request Forgery (CSRF) allows a third party to redirect static content within the security context of a trusted site.
Incorrect Answers and Explanations: B, C, and D. Answers B, C, and D are incorrect. Cross-Site Scripting (XSS) is third-party execution of web script-

ing languages such as Javascript within the security context of a trusted site. XSS is similar to CSRF; the difference is that XSS uses active code. PHP Remote File Inclusion (RFI) alters normal PHP variables to reference remote content, which can lead to execution of malicious PHP code. SQL Injection manipulates a back-end SQL server via a front-end web server.

4. Correct Answer and Explanation: **B**. Answer **B** is correct; Interface Definition Language (IDL) allows CORBA objects to communicate via a message interface.

Incorrect Answers and Explanations: **A**, **C**, and **D**. Answers **A**, **C**, and **D** are incorrect. Distributed Component Object Model (DCOM) is a Microsoft object broker that locates objects over a network. Object Linking and Embedding (OLE) is a part of DCOM that links documents to other documents. Object Management Guidelines is a distracter answer, playing off the term OMG: Object Management Group (OMG) developed CORBA.

5. Correct Answer and Explanation: **C**. Answer **C** is correct; database replication mirrors a live database, allowing simultaneous reads and writes by clients to multiple replicated databases.

Incorrect Answers and Explanations: **A**, **B**, and **D**. Answers **A**, **B**, and **D** are incorrect. DBMSs may attempt to commit updates—that is, make pending changes permanent. A database journal is a log of all database transactions. A shadow database is similar to a replicated database, with one key difference: A shadow database mirrors all changes made to a primary database, but clients do not access the shadow.

Endnotes

1. eXtreme Programming: A gentle introduction. URL: www.extremeprogramming.org/ (accessed July 23, 2010).
2. Generally Accepted Principles and Practices for Securing Information Technology Systems. URL: <http://csrc.nist.gov/publications/nistpubs/800-14/800-14.pdf> (accessed July 23, 2010).
3. Security-Related Policy Settings. URL: <http://technet.microsoft.com/en-us/library/bb457148.aspx> (accessed July 23, 2010).
4. CORBA® BASICS, URL: www.omg.org/gettingstarted/corbafaq.htm (accessed July 23, 2010).
5. McConnell, Steve. *Code Complete: A Practical Handbook of Software Construction*. Microsoft Press, 1993.
6. Watts Humphrey: *He Wrote the Book On Debugging*. URL: www.businessweek.com/magazine/content/05_19/b3932038_mz009.htm (accessed July 23, 2010).
7. 2010 CWE/SANS Top 25 Most Dangerous Software Errors. URL: <http://cwe.mitre.org/top25/> (accessed July 23, 2010).
8. Capability Maturity ModelSM for Software, Version 1.1. URL: www.sei.cmu.edu/reports/93tr024.pdf (accessed July 23, 2010).
9. Ibid.
10. The Database Normalization Process. URL: www.informit.com/articles/article.aspx?p=30646 (accessed July 23, 2010).

This page intentionally left blank

CHAPTER 9

Domain 9: Operations Security

147

Exam Objectives in this Chapter

- Administrative security
- Sensitive information/media security
- Continuity of operations
- Computer and network attacks

INTRODUCTION

The domain of operations security is concerned with threats to a production environment. Threat agents can be internal or external, and operations security must account for both in order to be effective. Ultimately operations security centers on the fact that people need appropriate access to data. This data exists on particular media, and is accessible by means of a system. Operations security is thus about people, data, media, and hardware, and the threats associated with them in a production environment.

ADMINISTRATIVE SECURITY

All organizations contain people, data, and the means for people to use the data. A fundamental aspect of operations security is ensuring that controls are in place to prevent people from either inadvertently or intentionally compromising the confidentiality, integrity, or availability of data or the systems and media holding it. Administrative security provides a way to control people's operational access to data.

Separation of duties

Separation of duties prescribes that more than one individual be required to complete critical or sensitive transactions. The goal is to ensure that, in order for someone to abuse his access to sensitive data or transactions, he must convince another party to act in concert. **Collusion** is the term used for two parties conspiring to undermine the security of the transaction. The classic action movie example of separation of duties involves two keys, a nuclear sub, and a rogue captain.

Rotation of duties/job rotation

By using rotation of duties, also known as job rotation or rotation of responsibilities, an organization can mitigate the risk associated with any one individual having too many privileges. Simply put, critical functions or responsibilities are not continuously performed by the same person without interruption. Rotation of duties can mitigate collusion.

Nondisclosure agreement

A nondisclosure agreement (NDA) is a work-related contract ensuring that, prior to being given access to sensitive information or data, an individual or organization appreciates the legal responsibility to maintain the confidentiality of sensitive information. Nondisclosure agreements are often signed by job candidates before they are hired, as well as by consultants or contractors. They are largely a directive control.

Background checks

Background checks (also known as background investigations or pre-employment screenings) are an additional administrative control commonly employed by many organizations. The majority of background investigations are performed as part of a pre-employment screening process. Some organizations conduct cursory background investigations that include a check of criminal records. Others do more in-depth checking, such as verifying employment history, obtaining credit reports, and in some cases requiring a drug test.

Configuration management

One of the most important components of any systems security work is developing a consistent system security configuration that can be leveraged throughout the organization. The goal is to move beyond the default system configuration to one that is hardened and that meets organizational requirements. One of the best ways to protect an environment from future zero day attacks (i.e., against vulnerabilities with no patch or fix) is to have a hardened system that provides only the functionality strictly required by the organization.

Basic configuration management practices in system security involve disabling unnecessary services; removing extraneous programs; enabling security capabilities such as firewalls, antivirus, and intrusion detection or prevention systems; and configuring security and audit logs.

Change management

A system that does not change becomes less secure over time, because security updates and patches are not applied. In order to maintain consistent and known operational security, a regimented change management, or change control, process needs to be followed. The purpose is to understand, communicate, and document any changes with the primary goal of understanding, controlling, and avoiding any direct or indirect negative impact that the change might impose.

Fast Facts

Because of the variability of the change management process, specific named phases have not been offered in this section. However, the general flow of change management includes

- Identifying a change
- Proposing a change
- Assessing the risk associated with the change
- Testing the change
- Scheduling the change
- Notifying impacted parties of the change
- Implementing the change
- Reporting results of the change

All changes must be closely tracked and auditable. A detailed change record should be kept. Some changes can destabilize systems or cause other problems; change management auditing allows operations staff to investigate recent changes in the event of an outage or problem. Audit records also allow verification that change management policies and procedures have been followed.

SENSITIVE INFORMATION/MEDIA SECURITY

Although security and controls related to personnel are vitally important, so is having a regimented process for handling sensitive information, including media. This section discusses concepts that are an important component of a strong overall information security program.

Storage

When storing sensitive data, it is preferable to encrypt it. Encryption of data at rest greatly reduces the likelihood of unauthorized disclosure due to media security issues. Physical storage of the media containing sensitive information should not be haphazard, whether the data is encrypted or not.

Media sanitization or destruction of data

It is time to destroy data or its associated media once an organization has identified that it is no longer operationally or legally required. While some data might not be sensitive and so not warrant thorough destruction, an organization will have data that must be verifiably destroyed or otherwise rendered nonusable if the media on which it is housed is recovered by a third party.

DATA REMANENCE

The term *data remanence* is important to understand when discussing media sanitization and data destruction. It refers to data that persists beyond noninvasive means to delete it. Though data remanence sometimes specifically refers to residual data that persists on magnetic storage, its concerns go beyond magnetic storage media. Security professionals must understand and appreciate the steps to make data unrecoverable.

WIPING OR OVERWRITING

When a user deletes a file, usually the file system merely removes metadata pointers or references to it. The file allocation table references are gone, but the file data remains. Significant amounts of “deleted data” may be recovered (“undeleted”) by readily available forensic tools. Reformatting a file system may also leave data intact.

Though simple file deletion or hard disk reformatting is not sufficient to render data unrecoverable, files may be securely wiped or overwritten. **Wiping**, also called *overwriting*, writes new data over each bit or block of file data. One of the shortcomings of wiping is that physical damage to a hard disk prevents complete overwriting.

DEGAUSSING

By introducing an external magnetic field with a **degausser**, the data on magnetic storage media can be made unrecoverable. Magnetic storage depends on the magnetization of the media being static unless intentionally changed by the storage device. A degausser destroys the integrity of the magnetization, making the data unrecoverable.

PHYSICAL DESTRUCTION

Physical destruction, when carried out properly, is considered the most secure means of media sanitization. One of the reasons for its higher degree of assurance is the greater likelihood of errors in data remanence with wiping or degaussing. Physical destruction is certainly warranted for the most sensitive data. Common means of destruction include incineration and pulverization.

SHREDDING

A simple form of media sanitization by physical destruction is shredding. Though sometimes used in relation to data overwriting, here *shredding* refers to the process of making data printed on hard copy, or on smaller objects such as floppy or optical disks, unrecoverable. Sensitive printed information needs to be shredded prior to disposal in order to thwart dumpster diving. This is a physical attack in which an attacker recovers trash in hopes of finding intact sensitive information.

CONTINUITY OF OPERATIONS

Although some continuity concepts were covered in Chapter 6, this section will focus on overtly operational continuity, which is principally concerned with availability.

Service Level Agreements

A Service Level Agreement (SLA) stipulates all expectations regarding the behavior of the department or organization that is responsible for providing services and the quality of those services. Often service level agreements dictate what is considered acceptable regarding bandwidth, time to delivery, response times, and so forth.

Fault tolerance

For systems and solutions within an organization to be able to provide continuous operational availability they must be implemented with fault tolerance in mind. Availability is not solely focused on system uptime requirements, but also requires that data be accessible in a timely fashion.

BACKUP

The most basic and obvious measure for increasing system or data fault tolerance is to provide for recoverability in event of failure. Given a long enough timeframe, accidents will happen. In order for data to be recoverable in case of a fault, some form of backup or redundancy must be in place. Although magnetic tape media are an old technology, they are still the most common repository for backup data. The three basic types of backup are **full**, **incremental**, and **differential**.

Full

A full backup is a replica of all allocated data on a hard disk. It contains all of the allocated data on the hard disk, which makes recovery simple. Because of the larger amount of media, and therefore the cost, and the longer backup window requirements, full backups are often coupled with either incremental or differential backups to balance the time and media considerations.

Incremental and Differential

Incremental backups archive only files that have changed since the last backup of any kind. Differential backups archive any files that have been changed since the last full backup.

DID YOU KNOW?

Assume that a full backup is performed every Sunday, and either incremental or differential backups are performed daily from Monday through Saturday. Data is lost after Wednesday's backup.

If incremental daily backups were carried out in addition to the weekly full backup, the tapes from Sunday, Monday, Tuesday, and Wednesday are needed to recover all archived data. If differential backups were carried out in addition to the full weekly backup, only the Sunday and Wednesday tapes are needed.

REDUNDANT ARRAY OF INEXPENSIVE DISKS

Even if only one full backup tape is needed for recovery of a system due to a hard disk failure, the time to recover a large amount of data can easily exceed the recovery time dictated by the organization. One goal of a Redundant Array of Inexpensive Disks (**RAID**) is to mitigate the risk of hard disk failures. The various RAID levels consist of different approaches to disk array configurations.

Fast Facts

Three terms are important to understand with respect to RAID: mirroring, striping, and parity.

Mirroring, the most obvious and basic RAID concept, is used to achieve full data redundancy by writing the same data to multiple hard disks. Since mirrored data must be written to multiple disks, the write times are slower. However, there can be performance gains when reading mirrored data from the disks simultaneously. Other than read and write performance considerations, a major cost associated with mirroring is disk usage; at least half of the drives are used for redundancy.

Striping focuses on increasing read and write performance by spreading data across multiple hard disks. This allows reads and writes to be performed in parallel across multiple disks rather than serially on one disk. Parallelization provides a performance increase, but does not enhance data redundancy.

Parity is a means to achieve data redundancy without incurring the cost of mirroring in terms of disk usage and write performance.

RAID 0—Striped Set

As suggested by the name, RAID 0 employs striping to increase the performance of reads and writes. By itself, striping offers no data redundancy, so RAID 0 is a poor choice if data recovery is the reason for leveraging RAID. Figure 9.1 shows RAID 0.

RAID 1—Mirrored Set

RAID 1 is perhaps the simplest RAID level to understand. It creates an exact duplicate of all data on an additional disk. Write performance is decreased although read performance can be increased. Disk cost is one of the most troubling aspects of this level because at least half of all disks are dedicated to redundancy. Figure 9.2 shows RAID 1.




FIGURE 9.1
RAID 0—Striped set.




FIGURE 9.2
RAID 1—Mirrored set.

RAID 2—Hamming Code

RAID 2 is not considered commercially viable for hard disks and is not used. It requires either 14 or 39 hard disks and a specially designed hardware controller, which makes it extremely cost prohibitive.

Exam Warning

While the ability to quickly recover from a disk failure is a goal of RAID, there are configurations that do not feature reliability. For the exam, be sure to understand that not all RAID configurations provide additional reliability.

RAID 3—Striped Set with Dedicated Parity (Byte Level)

Striping is desirable for the performance gains associated with spreading data across multiple disks. However, striping alone is not desirable because of the lack of redundancy. With RAID 3, data at the byte level is striped across multiple disks, but an additional disk is leveraged for storage of parity information, which is used for recovery from failure.

RAID 4—Striped Set with Dedicated Parity (Block Level)

RAID 4 provides exactly the configuration and functionality that RAID 3 offers, but stripes data at the block, rather than the byte, level. Like RAID 3, RAID 4 employs a dedicated parity drive rather than distributing parity data among all disks, as in RAID 5.

RAID 5—Striped Set with Distributed Parity

One of the most popular RAID configurations is RAID 5, Striped Set with Distributed Parity. RAID 5 focuses on striping for the performance increases it offers, and leverages block-level striping. Like RAID 3 and 4, it writes parity information that is used for recovery purposes. However, unlike RAID 3 and 4, which require a dedicated disk for parity information, RAID 5 distributes the parity information across multiple disks.

One of the reasons for RAID 5's popularity is that the disk cost for redundancy is lower than that associated with a mirrored set. Another important reason for its popularity is the support for both hardware- and software-based implementations, which significantly reduces the barrier to entry for RAID configurations. RAID 5 allows data recovery in the event that any one disk fails. Figure 9.3 illustrates RAID 5.

RAID 6—Striped Set with Dual Distributed Parity

While RAID 5 accommodates the loss of any one drive in the array, RAID 6 can withstand the failure of two drives and still function. This redundancy is achieved by writing the same parity information to two different disks.

RAID 1+0

RAID 1+0 is an example of nested RAID or multi-RAID, which simply means that one standard RAID level is encapsulated within another. With RAID 10,

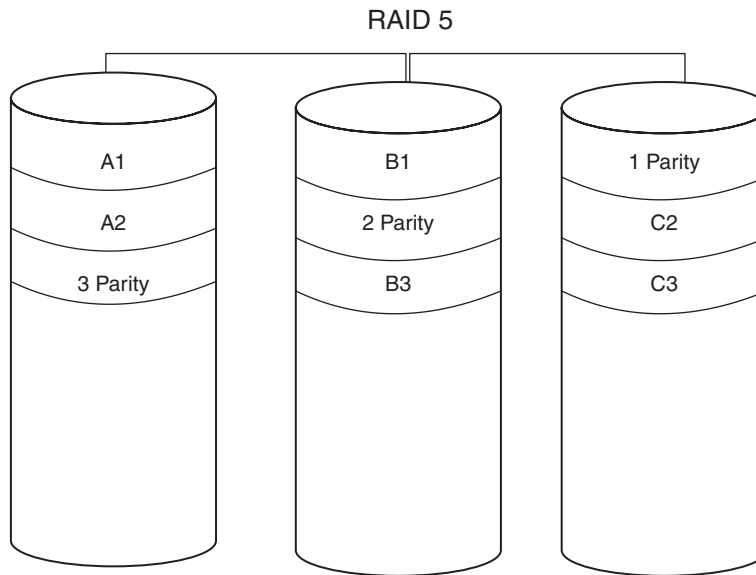


FIGURE 9.3
RAID 5—Striped Set with Distributed Parity.

which is also commonly referred to as RAID 1+0 to explicitly indicate the nesting, the configuration is that of a striped set of mirrors.

SYSTEM REDUNDANCY

The redundancy and resiliency of data provided by RAID and backup solutions is important. However, further consideration needs to be given to the systems themselves that provide access to this redundant data.

Crunch Time

Table 9.1 provides a brief description of the most common RAID levels.

Table 9.1 RAID Levels	
RAID Level	Description
RAID 0	Striped Set
RAID 1	Mirrored Set
RAID 3	Byte-Level Striping with Dedicated Parity
RAID 4	Block-Level Striping with Dedicated Parity
RAID 5	Block-Level Striping with Distributed Parity
RAID 6	Block-Level Striping with Double Distributed Parity

Redundant Hardware

Many systems can provide internal hardware redundancy of components that are prone to failure. The most common examples are systems or devices that have redundant onboard power in the event of a power supply failure. In addition to redundant power, it is common to find redundant network interface cards (NICs) as well as redundant disk controllers. Sometimes systems simply have field replaceable modular versions of commonly failing components. Physically replacing a power supply might increase downtime, but having an inventory of spare modules to repair the entire data center's servers is less expensive than having all servers configured with an installed redundant power supply.

Redundant Systems

Although quite a few fault-prone internal components can be configured for built-in system redundancy, internal redundancy does have its limits. If system availability is extremely important, it might be prudent to have entire systems available in inventory to serve as a means to recover. While the time to recover might be greater, it is fairly common for organizations to have an SLA with their hardware manufacturers to be able to quickly procure replacement equipment. If the recovery times are acceptable, quick procurement is likely to be far cheaper than having spare equipment on hand.

High-Availability Clusters

Extremely critical applications and systems have more stringent uptime requirements than can be met by standby redundant systems or spare hardware. These systems and applications typically require what is referred to as a high-availability (HA), or failover, cluster. A high-availability cluster employs multiple systems that are already installed, configured, and plugged in, so if a failure causes one of the systems to fail the other can be seamlessly leveraged to maintain the availability of the service or application being provided.

The primary implementation consideration for high-availability clusters is whether each node is actively processing data in advance of failure. This is known as an active-active configuration, and is commonly referred to as load balancing. Having systems in an active-active, or load-balancing, configuration is typically more costly than having them in an active-passive, or hot-standby, configuration, in which the backup systems begin processing only when a failure state is detected.

COMPUTER AND NETWORK ATTACKS

This section provides basic information on attacks commonly experienced by organizations. First, a brief discussion of threats, which will help to bring these attacks into the organizational risk assessment model.

Password guessing and password cracking

It is prudent to differentiate password guessing and password cracking, as the techniques differ. Password guessing, the simpler of the two from both the attacker's and the defender's vantage point, is an online technique for authenticating as a particular user to the system. Password cracking is an offline technique in which the attacker gains access to the password hashes or the database.

Password guessing may be detected by monitoring failed-login system logs. In order to differentiate between attackers and normal users accidentally mistyping their passwords, clipping levels are useful. **Clipping levels** define a minimum reporting threshold level.

Session hijacking and MITM

Session hijacking compromises an existing network session, sometimes seizing control of it. Older protocols such as Telnet may be vulnerable to session hijacking.

A Man in the Middle (MITM; also called Monkey in the Middle) attack places the attacker between the victim and another system: The attacker's goal is to be an undiscovered proxy for either or both of two endpoints engaging in communication. Effectively, an attacker suitably positioned through a combination of spoofing, masquerading as another endpoint, and sniffing traffic is potentially able to insert herself in the middle of a connection. The capabilities of session hijacking include changing content as it is delivered to one of the endpoints, initiating transactions as one side of the connection, distribution of malware to either end of the connection, and so on. Session hijacking can best be prevented by leveraging encrypted communications, which provide mutual endpoint authentication.

Malware

Malware, or malicious code/software, is one of the best known threats to information systems. Numerous types of malware, some detailed in Table 9.2, have evolved over the years to continually cause stress to operations.

Denial of Service and Distributed Denial of Service

Denial of Service (**DoS**) is a one-to-one availability attack; Distributed Denial of Service (**DDoS**) is a many-to-one availability attack. The two are among the easiest techniques to understand as they are simply directed against a site, system, or network. Though there are many local denial of service techniques, this section focuses on the remote variety. DoS attacks come in all shapes and sizes, ranging from those involving one specially crafted packet and a vulnerable system to see that packet, to DDoS attacks that leverage tens of thousands (or more) bots to target an online service provider with a flood of seemingly legitimate traffic in an attempt to overwhelm its capacity. Table 9.3 lists historical examples of malicious packet attacks as well as some general resource exhaustion, or flooding, techniques.

Table 9.2 Types of Malware

Malicious Code	Description
Virus	The term that most laypersons use for all bad things that can happen on a computer. Information security professionals require a bit more specificity and reserve <i>virus</i> to indicate malicious code that hooks onto executable code and requires user interaction to spread. In addition to spreading, the actual payload of the virus—that is, what it is intended to do—can be anything.
Macro virus	The term <i>macro virus</i> refers to malicious code that infects Microsoft Office documents by embedding malicious macros within them. Many organizations were wholly unaware of the macro functionality provided by Microsoft Office until being hit with macro viruses.
Worm	The distinguishing feature of <i>worms</i> is their ability to self-propagate, or, spread without user interaction. This has made them exceedingly good at spreading rapidly throughout the Internet. Some of the most well known malware falls under the worm category: Code Red, Nimda, SQL Slammer, Blaster, MyDoom, Witty.
Trojan horses	They get their name from the famous Trojan Horse of Greek mythology, are defined by how they are concealed, and are most often associated with providing an attacker with persistent backdoor access. Trojans ostensibly provide the functionality that the user is seeking, but also come with malicious functionality that the user does not anticipate.
Rootkit	The term <i>rootkit</i> is used for malware that focuses on hiding its existence from a savvy administrator trying to detect it. Typical capabilities include file, folder, process, and network connection hiding. The techniques developed with rootkits are now commonly included in other types of malware.

Table 9.3 Denial of Service Examples

DoS	Type	Description
Land	Malformed packet	The land attack uses a spoofed SYN packet that includes the victim's IP address as both source and destination. This attack targets the TCP/IP stack of older unpatched Windows systems.
Smurf	Resource exhaustion	A smurf attack involves ICMP flooding. The attacker sends ICMP Echo Request messages with spoofed source addresses of the victim to the directed broadcast address of a network known to be a smurf amplifier. A smurf amplifier is a public-facing network that is misconfigured such that it will forward packets sent to the network broadcast address to each host in the network.

Continued

Table 9.3 Denial of Service Examples—Cont'd

DoS	Type	Description
SYN Flood	Resource exhaustion	SYN floods are the most basic resource exhaustion attacks, and involve an attacker, or attacker-controlled machines, initiating many connections to the victim but not responding to the victim's SYN/ACK packets. The victim's connection queue will eventually be unable to process any more new connections.
Teardrop	Malformed packet	The teardrop attack involves a malformed packet that targets issues with systems' fragmentation reassembly. Packets are sent with overlapping fragment offsets, which can cause problems for a system attempting to reassemble the fragments.
Ping of Death	Malformed packet	The ping of death involves sending a malformed ICMP Echo Request (Ping) that is larger than the maximum size of an IP packet. Historically, sending the Ping of Death crashed systems.
Fraggle	Resource exhaustion	The fraggle attack is a variation of the smurf attack, the main difference being that fraggle leverages UDP for the request portion and stimulates, most likely, the sending of an ICMP Port Unreachable message to the victim rather than an ICMP Echo Response.
DNS Reflection	Resource exhaustion	A more recent technique that, like the smurf attack, leverages a third party. The attacker has poorly configured third-party DNS servers query an attacker-controlled DNS server and cache the response (a maximum-size DNS record). Once the large record is cached by many third-party DNS servers, the attacker sends DNS requests for those records with a spoofed source of the victim. This causes these extremely large DNS records to be sent to the victim in response.

SUMMARY OF EXAM OBJECTIVES

In this chapter we discussed operational security. This domain concerns the security of systems and data while they are being actively used in a production environment. Ultimately operations security is about people, data, media, and hardware, all of which are elements that need to be considered from a security perspective. The best technical security infrastructure in the world will be rendered useless if an individual with privileged access decides to turn against the organization and that organization has no preventive or detective controls in place to defend itself.

TOP FIVE TOUGHEST QUESTIONS

1. Which type of attack makes use of misconfigured third-party systems to perpetrate a denial of service?
 - A. Smurf
 - B. Session hijacking
 - C. Teardrop
 - D. Land
2. Which attack technique might involve a seemingly trusted endpoint resolving as a website hosting malware?
 - A. Password cracking
 - B. Trojan horse
 - C. Session hijacking
 - D. UI redressing
3. Which principle involves defining a trusted security baseline image of critical systems?
 - A. Configuration management
 - B. Change management
 - C. Patch management
 - D. Vulnerability management
4. Which type of attack leverages overlapping fragments to cause a denial of service?
 - A. Smurf
 - B. Teardrop
 - C. Fraggle
 - D. Session hijacking
5. Which security principle can be used to help detect fraud coming from users becoming comfortable in their position?
 - A. Separation of duties
 - B. Principle of least privilege
 - C. Rotation of duties
 - D. Collusion

Answers

1. Correct Answer and Explanation: A. Answer A is correct; smurf attacks are a DoS technique that uses spoofed ICMP Echo Requests sent to misconfigured third parties (amplifiers) in an attempt to exhaust the victim's resources.
Incorrect Answers and Explanations: B, C, and D. Answers B, C, and D are incorrect. Session hijacking involves a combination of sniffing and spoofing to allow the attacker to masquerade as one or both ends of an established connection. The teardrop attack works by sending overlapping fragments that, when received by a vulnerable host, can cause a system to crash. The land attack is a malformed packet DoS that can cause vulnerable systems to crash by sending a SYN packet with both the source and destination IP address set to that of the victim.

- 2. Correct Answer and Explanation: C.** Answer C is correct; session hijacking involves a combination of sniffing and spoofing so that the attacker can masquerade as one or both ends of an established connection.

Incorrect Answers and Explanations: **A, B, and D.** Answers **A, B, and D** are incorrect. Password cracking has little to do with which website is resolved. Though Trojan Horse infections no doubt have the ability to alter hosts tables, DNS settings, and other things that can cause this behavior, they are considered malware rather than an attack technique. Also the mention of a trusted endpoint makes session hijacking the more likely answer. UI redressing is a simple distraction answer, and is the more generic term for what is known as clickjacking.
- 3. Correct Answer and Explanation: A.** Answer A is correct; configuration management involves the creation of known security baselines for systems, which are often built leveraging third-party security configuration guides.

Incorrect Answers and Explanations: **B, C, and D.** Answers **B, C, and D** are incorrect. Change management is concerned with ensuring a regimented process for any system changes. Patch management focuses on ensuring that systems receive timely updates to the security and functionality of the installed software. The goal of vulnerability management is to understand what known vulnerabilities exist in an organization and to track their remediation over time.
- 4. Correct Answer and Explanation: B.** Answer B is correct; the teardrop attack is a DoS that works by sending overlapping fragments that, when received by a vulnerable host, can cause a system to crash.

Incorrect Answers and Explanations: **A, C, and D.** Answers **A, C, and D** are incorrect. Smurf attacks are a DoS that uses spoofed ICMP Echo Requests sent to misconfigured third parties (amplifiers) in an attempt to exhaust the victim's resources. Fraggle attacks are a smurf variation that uses spoofed UDP rather than ICMP messages to stimulate the misconfigured third-party systems. Session hijacking involves a combination of sniffing and spoofing in which the attacker masquerades as one or both ends of an established connection.
- 5. Correct Answer and Explanation: C.** Answer C is correct; rotation of duties is useful in detecting fraud by requiring that more than one employee perform a particular task. In addition to fraud detection, rotation can determine if there is a lack of depth for a given role or function within the organization.

Incorrect Answers and Explanations: **A, B, and D.** Answers **A, B, and D** are incorrect. Separation of duties attempts to prevent fraud by requiring multiple parties to carry out a transaction or by segregating conflicting roles. The principle of least privilege is not associated specifically with fraud detection. Collusion is the term for multiple parties acting together to perpetrate a fraud.

CHAPTER 10

Domain 10: Legal, Regulations, Investigations, and Compliance

161

Exam Objectives in this Chapter

- Major legal systems
- Criminal, civil, and administrative law
- Legal aspects of information security
- Legal aspects of investigations
- Important laws and regulations
- Ethics

INTRODUCTION

This chapter will introduce some of the basic legal concepts that are important to all information security professionals. The actual implementation of laws surrounding intellectual property, privacy, reasonable searches, and breach notification, to name a few, differ among various regions of the world, but the importance of these concepts is still universal.

MAJOR LEGAL SYSTEMS

In order to appreciate common legal concepts at work in today's global economy, an understanding of the major legal systems is required. These systems provide the framework that determines how a country develops laws pertaining to information systems in the first place. The three major systems of law are civil, common, and religious.

Civil law

By far the most common of the major legal systems is civil law, which is employed by many countries throughout the world. This system leverages codified laws or statutes to determine what is considered within legal bounds. Though a legislative branch typically wields the power to create laws, there still exists a judicial branch that must interpret them. The most significant difference between civil and common law is that, under civil law, judicial precedents and particular case rulings do not carry the weight they do under common law.

Common law

Common law is the legal system used in the United States, Canada, the United Kingdom, and most former British colonies, among others. Its primary distinguishing feature is the significant emphasis on particular cases and judicial precedent as a determinant of laws. Though typically a legislative body creates new statutes and laws, judicial rulings can, at times, supersede those laws. Because of the emphasis on judges' interpretations, there is significant possibility that as society changes over time, so too can the judicial interpretations.

Religious and customary law

Religious law serves as the third of the major legal systems. Religious doctrine or interpretation can be a source of legal understanding and statutes. However, the extent and degree to which religious texts, practices, or understandings are consulted can vary greatly.

Customary law refers to customs or practices that are so commonly accepted by a group that they are treated as a law. These practices can later be codified in the more traditional sense, but the emphasis on prevailing acceptance by a group is important with respect to the concept of negligence, which, in turn, is important in information security. The concept of "best practices" is closely associated with customary law.

CRIMINAL, CIVIL, AND ADMINISTRATIVE LAW

Common law is the most well represented system on the exam, so it is the primary focus in this chapter. Within common law there are various branches of laws, including criminal, civil, and administrative.

Criminal law

Criminal law pertains to laws where the victim can be seen as society itself. While it might seem odd to consider society the victim when an individual is murdered, the goal of criminal law is to promote and maintain an orderly and law abiding citizenry. Criminal law can include penalties that remove an individual from society by incarceration or, in some extreme cases in some regions, by death. The goals of criminal law are to deter crime and to punish offenders.

Because of the seriousness of potentially depriving someone of his freedom or, in the most extreme cases his life, the burden of proof in criminal cases is beyond any reasonable doubt.

Civil law

In addition to being a major legal system in the world, civil law also serves as a type of law within the common law legal system. Associated with civil law is tort law, which deals with injury, loosely defined, that results from an

Table 10.1 Common Types of Financial Damages

Financial damages	Description
Statutory	Statutory damages are those prescribed by law, which can be awarded to the victim even if she incurred no actual loss or injury.
Compensatory	The purpose of compensatory damages is to provide the victim with a financial award in an effort to compensate for the loss or injury incurred as a direct result of the wrongdoing.
Punitive	The intent of punitive damages is to punish an individual or organization. These damages are typically awarded in an attempt to discourage a particularly egregious violation where the compensatory or statutory damages alone would not act as a deterrent.

individual violating her responsibility to provide a duty of care. Tort law is the primary component of civil law and the most significant source of lawsuits seeking damages.

In the United States, the burden of proof in civil proceedings is the preponderance of the evidence. By “preponderance” is meant “more likely than not.” Satisfying the burden of proof requirement of the preponderance of the evidence in a civil matter is a much easier task than meeting the burden of proof requirement in criminal proceedings. The usual outcome of a successful ruling against a civil defendant is the payment of financial damages. The most common types of financial damages are presented in Table 10.1.

Administrative Law

Administrative or regulatory law is enacted by government agencies. In the United States, administrative law is enacted by the executive branch (deriving from the office of the President). Government-mandated compliance measures are administrative laws.

The executive branch can create administrative law without requiring input from the legislative branch, but the law must still operate within the confines of civil and criminal code, and can still come under scrutiny by the judicial branch. Some examples of administrative law are FCC, HIPAA, FDA, and FAA regulations.

INFORMATION SECURITY ASPECTS OF LAW

Examples of legal concepts that affect information security include crimes committed or aided by computer systems, attacks on intellectual property, privacy concerns, and international issues.

Computer Crime

One aspect of the interaction between information security and the legal system is represented by computer crimes. Applicable computer crime laws vary throughout the world, according to jurisdiction. However, some generalities do exist.

Fast Facts

Computer crimes can be understood based upon the way in which computer systems relate to the wrongdoing: as targets or as tools to perpetrate a crime.

- *Computer system as target:* Crimes where computer systems are a primary target, such as disrupting online commerce by means of Distributed Denial of Service attacks, installing malware on systems for the distribution of spam, or exploiting a system vulnerability to store illegal content.
- *Computer as tool:* Crimes where the computer is a central component in the commission of a crime. Examples include stealing trade secrets by compromising a database server, leveraging computers to steal cardholder data from payment systems, conducting computer-based reconnaissance to target an individual for information disclosure or espionage, and harassment.

International Cooperation

To date, the most significant progress toward international cooperation in computer crime policy is the Council of Europe Convention on Cybercrime, which a majority of the 47 European member countries, and the United States, have signed and ratified. The primary focus of the Convention on Cybercrime is establishing policy standards to promote international cooperation during investigation and prosecution of cybercrime. Additional information on the Council of Europe Convention on Cybercrime can be found at <http://conventions.coe.int/Treaty/en/Treaties/Html/185.htm>.

Due care and due diligence

The standard of **due care**, or duty of care, provides a framework for defining a minimum standard of protection that business stakeholders must achieve. Due care discussions often reference the Prudent Man Rule, and require that organizations engage in business practices that a prudent, right-thinking person would consider appropriate. Businesses that are found not to have applied this minimum duty of care can be deemed negligent in carrying out their duties.

Due diligence requires that an organization continually scrutinize its own, and any ancillary organization's, practices to ensure that they meet or exceed the requirements to protect assets and stakeholders. Due diligence is the management of due care: It follows a formal process.

Intellectual property

As opposed to physical or tangible property, intellectual property is intangible—the result of a creative act. The purpose of intellectual property law is to control the use of intangible property that can often be trivial to reproduce or abuse

once made public or known. The following concepts effectively create an exclusive monopoly on the use of intellectual property.

TRADEMARK

Trademarks are associated with marketing: Their purpose is to allow the creation of a brand that distinguishes the source of products or services. A distinguishing name, logo, symbol, or image represents the most commonly trademarked items. In the United States, two different symbols are used: The [™] symbol, which can be used freely to indicate an unregistered mark, is shown in Figure 10.1. The ® symbol, which is used with marks that have been formally registered with the U.S. Patent and Trademark Office, is shown in Figure 10.2.

Syngress™

FIGURE 10.1
Trademark symbol.

Syngress®

FIGURE 10.2
Registered trademark symbol.

PATENT

Patents give a monopoly to the patent holder on using, making, or selling an invention for a period of time, in exchange for the patent holder's making the invention public. During the life of the patent, the patent holder can, through civil litigation, exclude others from leveraging his invention. The length of time that a patent is valid (the patent term) varies by country, and by type of invention. Generally, in both Europe and the United States the patent term is 20 years from the initial filing date.

COPYRIGHT

Copyright protects the form of expression in artistic, musical, or literary works, and is typically denoted by the © symbol, as shown in Figure 10.3. The purpose is to preclude unauthorized duplication, distribution, or modification. Note that the form of expression is protected rather than the subject matter or ideas expressed.

© 2010 Syngress

FIGURE 10.3
Copyright symbol.

LICENSES

Software licenses are a contract between a vendor and a consumer. Though there are licenses that explicitly permit the consumer to do virtually anything with the software, including modify it for use in another commercial product, most commercial software licensing explicitly limits use and distribution. End-user license agreements (EULAs) are unusual because use of the software typically constitutes contractual agreement.

TRADE SECRETS

Trade secrets are business-proprietary information that is important to an organization's ability to compete, which the organization must exercise due care and due diligence to protect. Some of the most common protection methods used are noncompete and nondisclosure agreements (NDAs).

LEGAL ASPECTS OF INVESTIGATIONS

Investigations are a critical way in which information security professionals come into contact with the law. They are commonly carried out by forensics and incident response personnel, both of whom need a basic understanding of legal

matters to ensure that the legal merits of the investigation are not unintentionally tarnished. Evidence, and the appropriate method for handling it, is a critical legal issue that all information security professionals must understand.

Digital Forensics

Digital forensics provides a formal approach to investigations and evidence with special consideration of their legal aspects. The forensics process must preserve the “crime scene” and the evidence taken from it in order to prevent unintentional violation of the integrity of either the data or the data’s environment. A primary goal of forensics is to prevent unintentional system modification. Live forensics includes taking a bit-by-bit, or binary, image of physical memory, gathering details about running processes, and gathering network connection data.

In addition to the valuable data gathered during live forensic capture, the main source of forensic data typically comes from binary images of secondary and portable storage devices such as hard disk drives, USB flash drives, CDs, DVDs, and possibly associated cellular phones and mp3 players.

Normal backup software will only capture the active partitions of a disk, and, further, only data marked as allocated. Normal backups may miss forensically significant data, so binary images are used.

Fast Facts

To understand the difference between a binary image and a normal backup, the investigator needs to understand the four types of data.

- *Allocated space*: Portions of a disk partition marked as actively containing data.
- *Unallocated space*: Portions of a disk partition that do not contain active data. They include memory that has never been allocated and previously allocated memory that has been marked as unallocated. If a file is deleted, the portions of the disk that held the deleted file are marked as unallocated and available for use.
- *Slack space*: Data is stored in specific size chunks known as clusters. A cluster is the minimum size that can be allocated by a file system. If a particular file, or final portion of a file, does not require the entire cluster, some extra space exists within it. This is known as slack space: It may contain old data, or it can be used intentionally by attackers to hide information.
- *“Bad” blocks/clusters/sectors*: Hard disks routinely end up with sectors that cannot be read because of some physical defect. Since no data can be read in them, the sectors marked as bad are ignored by the operating system. Attackers can intentionally mark sectors or clusters as bad in order to hide data within them.

Evidence

Evidence is one of the most important legal concepts for information security professionals to understand. Information security professionals are commonly involved in investigations, and so often have to obtain or handle evidence.

Crunch Time

Real evidence consists of tangible and/or physical objects. A knife or a bloody glove might constitute real evidence in some traditional criminal proceedings. Direct evidence is testimony provided by a witness regarding what the person actually experienced with her five senses. Circumstantial evidence is

evidence that establishes the circumstances related to particular points or even other evidence. Hearsay evidence is second-hand; as opposed to direct evidence, it has not been directly seen or heard by the witness.

Evidence integrity

Evidence must be reliable. Because it is common during forensic and incident response investigations to analyze digital media, the integrity of the data must be maintained during the course of its acquisition and analysis. Checksums can ensure that no data changes occurred as a result of acquisition and analysis. One-way hash functions such as MD5 or SHA-1 are commonly used for this purpose. Chain of custody requires that once evidence is acquired, full documentation regarding who, what, when, and where it was handled is maintained. Figure 10.4 shows a chain of custody form.

IMPORTANT LAWS AND REGULATIONS

An entire book could easily be filled with discussions of both U.S. and international laws that directly or indirectly pertain to information security. This section is not an exhaustive review of these laws. Instead only those that are represented on the examination will be discussed.

Computer Fraud and Abuse Act

Title 18 of the United States Code Section 1030, which is more commonly known as the Computer Fraud and Abuse Act, was originally drafted in 1984 but is still an important piece of legislation related to the prosecution of computer crimes. The law has been amended numerous times, most notably by the PATRIOT Act of 2001 and by the more recent Identity Theft Enforcement and Restitution Act of 2008, which is too new to be included on the exam at the time of this writing.

The goal of the Computer Fraud and Abuse Act is to deter and prosecute acts that damage federal interest computers. "Federal interest" includes government, critical infrastructure, and financial processing systems; it also references computers used in interstate commerce. With the ubiquity of Internet-based commerce, this definition can justify almost any Internet-connected computer as being protected. The Computer Fraud and Abuse Act criminalizes actions involving intentional attacks against protected computers that result in aggregate damages of \$5,000 in one year.

CHAIN OF CUSTODY REPORT

DATE OF INCIDENT: _____

TIME: _____

NAME OF EMPLOYEE: _____

EMPLOYEE NUMBER: _____

DEPARTMENT: _____

IMMEDIATE SUPERVISOR: _____

INCIDENT DESCRIPTION: _____

DESCRIPTION OF EVIDENCE: _____

PERSON SEIZING EVIDENCE: _____ INITIALS: _____

LOCATION OF EVIDENCE WHEN SEIZED: _____

WITNESSES: _____

LAW ENFORCEMENT OFFICER: _____ INITIALS: _____

DATE AND TIME EVIDENCE TRANSFER: _____

FIGURE 10.4
Chain of custody form.

PATRIOT Act

The PATRIOT Act of 2001 was passed in response to the attacks on the United States that took place on September 11, 2001. Its full title is "Uniting and Strengthening America by Providing Appropriate Tools Required to Intercept and Obstruct Terrorism." The main thrust of the Act, as it applies to information security professionals, is less stringent oversight of law enforcement regarding data collection. Wiretaps have become broader in scope. Searches and seizures can be carried out without immediate notification of the person whose data or property might be seized. Additionally, the Act amends the Computer Fraud and Abuse Act to strengthen penalties for those convicted of attempting to damage a protected computer such that conviction of a second offense can mean up to 20 years in prison.

Privacy

One of the unfortunate side effects of the explosion of information systems over the past few decades is the loss of privacy. As more and more data about individuals is used and stored by information systems, the likelihood of it being inadvertently disclosed, sold to a third party, or intentionally compromised by a malicious insider or third party increases.

PRIVACY ACT OF 1974

All governments have a wealth of personally identifiable information about their citizens. The Privacy Act of 1974 was created to codify protections of U.S. citizens' data that is used by the federal government. It defines guidelines regarding how citizens' personally identifiable information can be used, collected, and distributed. An additional protection allows individuals to have access to the data related to them, limited only by some national security-oriented exceptions.

EUROPEAN UNION PRIVACY

The European Union has taken an aggressive pro-privacy stance while balancing the needs of business. Commerce would be impacted if member nations had different regulations regarding the collection and use of personally identifiable information. Therefore, the EU Data Protection Directive allows the free flow of information tempered by consistent protections of the data belonging to the citizens of each member nation.

Fast Facts

The principles of the EU Data Protection Directive are

- Notifying individuals regarding how their personal data is collected and used
- Allowing individuals to opt out of sharing their personal data with third parties
- Requiring individuals to opt in to sharing their most sensitive personal data
- Providing reasonable protections for personal data

OECD PRIVACY GUIDELINES

The Organization for Economic Cooperation and Development (OECD), although often considered exclusively European, consists of 30 member nations from throughout the world. In addition to prominent European countries, those members include such countries as the United States, Mexico, Australia, Japan, and the Czech Republic. The OECD is a forum for discussion of issues that impact the global economy. It routinely issues consensus recommendations that can serve as an impetus to changes in current policy and legislation in the OECD member countries and beyond.

EU-U.S. SAFE HARBOR

An interesting aspect of the EU Data Protection Directive is that the personal data of EU citizens may not be transmitted, even when permitted by the individual, to countries beyond the EU unless the receiving country is perceived by the EU

to have adequate data protection laws. This presents a challenge regarding the sharing of data with the United States, which is perceived to have less stringent privacy protections. To resolve this issue, the United States and the European Union created the safe harbor framework to give U.S. organizations the benefit of authorized data sharing. To be part of the Safe Harbor, U.S. organizations must voluntarily consent to data privacy principles that are consistent with the EU Data Protection Directive.

ETHICS

Though ethical behavior is desirable in all fields, the sensitive nature of information security makes an appreciation for ethics necessary for all information security professionals. Even more than in other fields, questionable ethical behavior in information security can be a career-ending move.

The study of ethics dates back thousands of years, and there is certainly no shortage of ethical paradigms from which to choose. This section will focus on the three ethical guidelines most likely to be covered by the CISSP exam.

Exam Warning

Although this chapter covers the Computer Ethics Institute and IAB codes of ethics, The (ISC)² © Code of Ethics is, understandably, far more important than either of those for examination purposes.

Computer Ethics Institute

The Computer Ethics Institute offers its **Ten Commandments of Computer Ethics** as a code of behavior for security professionals. This code is short and fairly straightforward. Both the name and format are reminiscent of the Ten Commandments of Judaism, Christianity, and Islam, but there is nothing overtly religious about them. The Computer Ethics Institute's Ten Commandments of Computer Ethics are as follows¹:

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

IAB's Ethics and the Internet

Much like the fundamental protocols of the Internet, the code of ethics of the Internet Activities Board (IAB), is defined in an RFC document: RFC 1087, Ethics and the Internet, published in 1987. According to the IAB, the following practices are considered unethical behavior²:

- Seeking to gain unauthorized access to the resources of the Internet.
- Disrupting the intended use of the Internet.
- Wasting resources (people, capacity, computer) through such actions.
- Destroying the integrity of computer-based information.
- Compromising the privacy of users.

The (ISC)² © Code of Ethics

The (ISC)² © Code of Ethics consists of a preamble as well as four canons, adherence to which is considered mandatory for certification as a CISSP. Both the preamble and the canons are printed on each CISSP certification card. The full (ISC)² © Code of Ethics is available at www.isc2.org/ethics/default.aspx.

DID YOU KNOW?

The four canons of the (ISC)² © Code of Ethics are specifically ordered according to their importance. Should a security professional ever find herself in a situation where the canons are at odds with one another, priority should be given based on the order in which the canons appear in the Code.

The first, and therefore most important, canon requires the information security professional to “protect society, the commonwealth, and the infrastructure.”³ The focus of the first canon is on the public and its understanding of and faith in information systems. Security professionals are charged with promoting safe security practices and bettering systems and infrastructure security for the public good.

The second canon charges information security professionals to “act honorably, honestly, justly, responsibly, and legally.”⁴ This is fairly straightforward, but there are a few points worth emphasizing. One is related to laws from different jurisdictions that conflict. The (ISC)² © Code of Ethics suggest that priority be given to the jurisdiction in which services are provided. Another point made by this canon is related to providing prudent advice and cautioning the security professional to avoid unnecessarily promoting fear, uncertainty, and doubt.

The third canon requires that security professionals “provide diligent and competent service to principals.”⁵ The primary focus here is ensuring that the security professional provide quality service for which she is qualified and which maintains the value and confidentiality of information and the associated systems. An additional important consideration is ensuring that the professional does not have a conflict of interest in providing quality services.

The fourth and final canon mandates that information security professionals “advance and protect the profession.”⁶ This requires that the security professional maintain his skills and advance the skills and knowledge of others. An additional requirement is that individuals do not negatively impact the security profession by associating professionally with those who might harm it.

SUMMARY OF EXAM OBJECTIVES

Understanding and appreciating legal systems, concepts, and terms are required of an information security practitioner working in the information-centric world of today. Maintaining the integrity of evidence, using hashing algorithms for digital evidence, and maintaining a provable chain of custody are vital.

The nature of information security and its inherent sensitivity make ethical frameworks an additional point requiring attention. This chapter presented the IAB’s RFC on Ethics and the Internet, the Computer Ethics Institute’s Ten Commandments of Computer Ethics, and The (ISC)² © Code of Ethics. The CISSP exam will, no doubt, emphasize the Code of Ethics proffered by (ISC)², which presents an ordered set of four canons that attend to responsibility to the public, individual behavior, competent service, and the profession as a whole.

TOP FIVE TOUGHEST QUESTIONS

1. Without the _____ or some other separate agreement, the EU Data Protection Directive would challenge the sharing of data with U.S. organizations because of the perceived lesser concern for privacy in the United States.
 - A. US-EU Safe Harbor
 - B. EU Privacy Harbor doctrine
 - C. Identity Theft Enforcement and Restitution Act
 - D. Federal Privacy Act
2. What can be used to make an exact replica of a hard disk drive as part of the evidence acquisition process?
 - A. Disk imaging software
 - B. Partition archival tool
 - C. Binary backup utility
 - D. Memory dumper
3. Which of the following defines “protected computers” and criminalizes attacks against them?
 - A. PATRIOT Act
 - B. Computer Fraud and Abuse Act
 - C. ECPA
 - D. Identity Theft Enforcement and Restitution Act
4. Which canon of The (ISC)² © Code of Ethics should be considered the most important?
 - A. Protect society, the commonwealth, and the infrastructure
 - B. Advance and protect the profession

- C.** Act honorably, honestly, justly, responsibly, and legally
 - D.** Provide diligent and competent service to principals
5. Which principle requires that an organization's stakeholders act prudently in ensuring that the minimum safeguards are applied to the protection of corporate assets?
- A.** Due protection
 - B.** Due process
 - C.** Due diligence
 - D.** Due care

Answers

1. Correct Answer and Explanation: **A.** Answer **A** is correct; the US-EU Safe Harbor agreement provides a framework in which U.S. companies can be considered safe for EU states and companies to share data with.

Incorrect Answers and Explanations: **B, C, and D.** Answers **B, C, and D** are incorrect. The EU Privacy Harbor doctrine is a made up choice. The other two options are legitimate U.S. laws important to information security, but neither specifically addresses the issues regarding data sharing with the EU.

2. Correct Answer and Explanation: **C.** Answer **C** is correct; a binary backup utility is needed to ensure that every single bit on a hard drive is copied. Slack and unallocated space is needed for a forensically sound image.

Incorrect Answers and Explanations: **A, B, and D.** Answers **A, B, and D** are incorrect. The most viable, but incorrect, choice is **A**, disk imaging software. While some disk imaging software provides bit-by-bit backup capabilities, typical usage only acquires allocated space. **D**, memory dumper, applies to physical memory, not a hard disk drive. **B** is a made-up phrase that sounds legitimate.

3. Correct Answer and Explanation: **B.** Answer **B** is correct; the Computer Fraud and Abuse Act, penned in 1984, is still an important piece of legislation for the prosecution of computer crime. It defines protected computers, which are those in which the federal government has a particular interest. The law sets a bar of \$5,000 in damages during one year for an action to constitute a crime.

Incorrect Answers and Explanations: **A, C, and D.** Answers **A, C, and D** are incorrect. The PATRIOT Act lessens some of the restrictions on law enforcement related to electronic monitoring. ECPA is concerned with the wiretapping of electronic communications. The Identity Theft Enforcement and Restitution Act of 2008 amends the Computer Fraud and Abuse Act to bring some of the considerations up to date.

4. Correct Answer and Explanation: **A.** Answer **A** is correct; to protect society, the commonwealth, and the infrastructure is the first canon, and is thus the most important of the four canons of The (ISC)² © Code of Ethics.

Incorrect Answers and Explanations: **B, C, and D.** Answers **B, C, and D** are incorrect. The canons of The (ISC)² © Code of Ethics are presented in order of importance. The second canon requires the security professional to act honorably, honestly, justly, responsibly, and legally. The third mandates that

professionals provide diligent and competent service to principals. The final, and therefore least important, canon requires professionals to advance and protect the profession.

5. Correct Answer and Explanation: **D.** Answer **D** is correct; due care provides a minimum standard of care that must be met. There are no explicit requirements that define due care. Rather, due care requires acting in accord with what a prudent person would consider reasonable.

Incorrect Answers and Explanations: **A, B, and C.** Answers **A, B, and C** are incorrect. Due protection is a made-up phrase that has no legal standing. Due process ensures that defendants are treated fairly in legal proceedings with respect to their constitutional rights. Due diligence is most closely related to the correct answer, due care. However, due diligence focuses on continually investigating business practices to ensure that due care is maintained.

Endnotes

1. Computer Ethics Institute. Ten Commandments of Computer Ethics. 1992. URL: <http://computerethicsinstitute.org/publications/tencommandments.html> (accessed September 13, 2010).
2. Internet Activities Board. RFC 1087—Ethics and the Internet. 1989. URL: <http://tools.ietf.org/html/rfc1087> (accessed September 13, 2010).
3. (ISC)² © Code of Ethics. No date. URL: www.isc2.org/ethics/default.aspx (accessed September 13, 2010).
4. Ibid.
5. Ibid.
6. Ibid.

802.1X Port Based Network Access Control; layer 2 authentication.

Abstraction The concealment of unnecessary details from the user.

Accountability An individual's responsibility for his or her actions.

Accreditation The Data Owner's acceptance of the risk represented by a system.

AES Advanced Encryption Standard; a block cipher using 128-bit, 192-bit, or 256-bit keys to encrypt 128-bit blocks of data.

Aggregation A mathematical attack where a user is able to use lower-level access to learn restricted information.

AH Authentication Header; an IPsec protocol that provides authentication and integrity for each packet of network data.

ALE Annualized Loss Expectancy; the cost of loss due to a risk over one year.

ALU Arithmetic Logic Unit; a CPU component that performs mathematical calculations.

Applet A small piece of mobile code that is embedded in other software such as a web browser.

ARO Annual Rate of Occurrence; the number of losses suffered per year.

Assembly language A low-level computer programming language with instructions that are short mnemonics, such as ADD, SUB (subtract), and JMP (jump), that match to machine language instructions.

Asymmetric Encryption Encryption that uses two keys: one to encrypt; the other, to decrypt.

ATM Asynchronous Transfer Mode; a WAN technology that uses fixed-length cells.

Attribute A column in a relational database table.

Authentication Proof of an identity claim.

Authorization Actions an individual is allowed to perform on a system.

Availability The assurance that information is available when needed.

Backdoor A shortcut in a system that allows a user to bypass security checks.

Baseline A uniform way to implement a safeguard; an administrative control.

Bell-LaPadula A security model focused on maintaining the confidentiality of objects.

BIA Business Impact Analysis; an analysis to identify and prioritize critical IT systems and components.

Biba A security model focused on maintaining the integrity of objects.

BIOS Basic Input Output System.

Bollard A post designed to stop a car, typically deployed in front of building entrances.

BRP Business Recovery Plan; the steps required to restore normal business operations after a disruptive event. Also known as *Business Resumption Plan*.

CCMP Counter Mode CBC; a MAC Protocol used by WPA2 to create a Message Integrity Check (MIC).

CCTV Closed Circuit Television; a television system used to detect the presence of intruders in restricted areas.

Certification A detailed inspection that verifies that a system meets documented security requirements.

CHAP Challenge Handshake Authentication Protocol; a secure network authentication protocol that uses a shared secret.

Cipher A cryptographic algorithm.

Ciphertext An encrypted message.

CISC Complex Instruction Set Computer; CPU instructions that are longer and more powerful than RISC.

Clipping level A minimum reporting threshold level.

Closed system A system that uses proprietary hardware or software.

- CMM** Capability Maturity Model; a maturity framework for evaluating and improving the software development process.
- CMP** Crisis Management Plan.
- COBIT** Control Objectives for Information and related Technology; a control framework for employing information security governance best practices within an organization.
- Collusion** An agreement between two or more individuals to subvert the security of a system.
- COM** Component Object Model; a standard that locates and connects objects locally.
- Compiler** The means by which source code, such as C or Basic, is converted into machine code.
- Confidentiality** The prevention of unauthorized disclosure of information.
- Constrained user interface** An interface that presents a user with limited controls on information, such as an ATM keypad.
- Continuity of Support Plan** A recovery plan that focuses narrowly on support of specific IT systems and applications.
- Control unit** A CPU component that acts as a traffic controller, sending instructions to the ALU.
- COOP** Continuity of Operations Plan; a plan to maintain operations during a disaster.
- Copyright** A type of intellectual property protection that protects the form of expression in artistic, musical, or literary works.
- CORBA** Common Object Request Broker Architecture; an open vendor-neutral networked object broker framework.
- Covert channel** Any communication that violates security policy.
- CPU** Central Processing Unit; the “brains” of the computer, capable of controlling and performing mathematical calculations.
- Crippleware** Partially functioning proprietary software, often with key features disabled. The user is typically required to make a payment to unlock the full functionality.
- Cryptanalysis** The science of breaking encrypted messages (i.e., recovering their meaning).
- Cryptography** The science of creating messages whose meaning is hidden.
- Cryptology** The science of secure communications.
- DAC** Discretionary Access Control, giving subjects full control of objects they have, or have been given, access to, including sharing with other subjects.
- DAD** Disclosure, Alteration, and Destruction; this is the opposite of Confidentiality, Integrity, and Availability.
- Data mining** A means of uncovering patterns, such as fraudulent activity, in a data warehouse.
- DBMS** Database Management System; the main control of all access to the database and the enforcer of database security.
- DCOM** Distributed Component Object Model; locates and connects objects across a network.
- DDL** Data Definition Language; a language for creating, modifying, and deleting tables.
- DDoS** Distributed Denial of Service; an availability attack using many systems.
- Decryption** Conversion of ciphertext into plaintext.
- Defense in depth** The application of multiple safeguards that span multiple domains to protect an asset.
- Degaussing** The destruction of the magnetization of storage media, thus making the data unrecoverable.
- Demarc** The demarcation point where the ISP’s responsibility ends and the customer’s begins.
- DES** Data Encryption Standard; a symmetric block cipher using a 56-bit key and a 64-bit block size.
- Diameter** The successor to RADIUS, designed to provide an improved Authentication, Authorization, and Accounting (AAA) framework.
- Differential backup** An archive of any files that have been changed since the last full backup was performed.
- Diskless workstation** A computer system that contains CPU, memory, and firmware, but no hard drive; a type of thin client.
- DML** Data Manipulation Language; a language for querying and updating data stored in tables.
- DNS** Domain Name System; a distributed global hierarchical database that translates names into IP addresses, and vice versa.
- DoS** Denial of Service; an attack on availability.

- DRAM** Dynamic Random Access Memory; storage of bits in small capacitors (e.g., small batteries) that is cheaper and slower than SRAM.
- DSL** Digital Subscriber Line; a system that uses existing copper pairs to provide digital service to homes and small offices.
- DSSS** Direct Sequence Spread Spectrum; a modulation technique that uses the entire wireless band at once.
- Due care** A requirement that key organizational stakeholders be prudent in carrying out their duties; also known as the "Prudent Man Rule."
- Due diligence** The management of due care.
- Dumpster diving** A physical attack in which the attacker recovers trash in hopes of finding sensitive information that has been merely discarded rather than destroyed.
- E1** A dedicated 2.048-megabit circuit that carries 30 channels.
- E3** The equivalent of 24 E1s.
- EAP** Extensible Authentication Protocol; a Layer 2 authentication framework that describes many specific authentication protocols.
- EAP-TLS** EAP-Transport Layer Security; an authentication framework that uses Public Key Infrastructure, requiring both server-side and client-side certificates.
- EAP-TTLS** EAP-Tunneled Transport Layer Security; a simplification of EAP-TLS that drops the client-side certificate requirement.
- EEPROM** Electrically Erasable Programmable Read Only Memory; memory that can be electrically erased via a flashing program.
- EF** Exposure Factor; the percentage of an asset's value lost due to an incident.
- Encryption** The conversion of plaintext to ciphertext.
- EPROM** Erasable Programmable Read Only Memory; memory that can be erased with ultraviolet light.
- ESP** Encapsulating Security Payload; an IPsec protocol which primarily provides confidentiality by encrypting packet data.
- Ethernet** The dominant local area networking technology, which transmits network data via frames.
- FHSS** Frequency Hopping Spread Spectrum; a modulation technique that uses a number of small frequency channels throughout the wireless band and "hops" through them in pseudorandom order.
- Free software** A controversial term that is defined differently by different groups. "Free" may mean free of charge, or it may mean that the user is free to use the software in any way he chooses, including modifying it.
- Freeware** Software that is free of charge.
- FTP** File Transfer Protocol; a method for transferring files to and from servers.
- Full backup** An archive of all files.
- GAN** Global Area Network; a global collection of WANs.
- Guideline** A recommendation; administrative control.
- Hash Function** One-way encryption using an algorithm and no key.
- Hearsay** Second-hand evidence not directly experienced by a witness.
- HMAC** Hashed Message Authentication Code; authentication that provides integrity by combining symmetric encryption with hashing.
- HTML** HyperText Markup Language; a language used to display web content.
- HTTP** HyperText Transfer Protocol; a protocol to transmit web data via a network.
- HTTPS** HyperText Transfer Protocol Secure; HTTP using SSL or TLS.
- ICC** See *Smartcard*.
- ICMP** Internet Control Message Protocol, a helper protocol that helps Layer 3.
- Identification** Association of an individual to a system
- IDL** Interface Definition Language; the language used by CORBA objects to communicate.
- IKE** Internet Key Exchange; a protocol that manages the IPsec encryption algorithm.
- IMAP** Internet Message Access Protocol; an email client protocol.

- Impact** The severity of damage, sometimes expressed in dollars (value).
- Incremental backup** An archive of all files that have changed since the last backup of any kind was performed.
- Inference** A deductive attack where the attacker uses lower-level access in order to learn restricted information.
- Integrity** Prevention of unauthorized modification of information.
- Interpreted code** Code that is compiled on the fly each time a program is run.
- Interrupt** An indication that an asynchronous CPU event has occurred.
- IPsec** Internet Protocol Security; a suite of protocols that provide a cryptographic layer to both IPv4 and IPv6.
- IPv4** Internet Protocol version 4, commonly called IP; the fundamental protocol of the Internet.
- IPv6** Internet Protocol version 6; the successor to IPv4, featuring a far larger address space, simpler routing, and simpler address assignment.
- ISAKMP** Internet Security Association and Key Management Protocol; the protocol that manages the IPsec Security Association process.
- ISDN** Integrated Services Digital Network; a set of standards for providing digital service via copper pair.
- ITIL** Information Technology Infrastructure Library; a framework for providing best services in IT Service Management.
- ITSEC** Information Technology Security Evaluation Criteria; the first successful international evaluation model.
- Kernel** The heart of the operating system, which usually runs in ring 0, providing the interface between hardware and the rest of the operating system, including applications.
- LAN** Local Area Network; a comparatively small network, typically confined to a building or an area within one.
- Lattice-Based Access Control** A nondiscretionary access control with defined upper and lower bounds implemented by the system.
- Layering** The division of hardware and software functionality into modular tiers.
- LEAP** Lightweight Extensible Authentication Protocol; a Cisco-proprietary protocol released before 802.1X was finalized.
- Least privilege** See *Principle of least privilege*.
- Logic bomb** A malicious program that is triggered when a logical condition is met, such as after a number of transactions have been processed or on a specific date.
- Lumen** The amount of light one candle creates.
- Lux** One lumen per square meter.
- MAC (Access Control)** Mandatory Access Control; a system-enforced access control based on a subject's clearances and an object's labels.
- MAC (Telecommunications)** Media Access Control; a Layer 2 protocol that transfers data to and from the physical layer.
- Machine code** Software that is executed directly by the CPU.
- Maintenance hook** A shortcut installed by system designers and programmers to allow developers to bypass normal system checks during development.
- Malware** Malicious software; any type of software that attacks an application or system.
- MAN** Metropolitan Area Network; a network typically confined to a city, a zip code, or a campus or office park.
- Mantrap** A preventive physical control with two doors. Each door requires a separate form of authentication to be opened.
- MD5** A hash function that creates a 128-bit message digest.
- MIC** Message Integrity Check; an integrity protocol used by WPA2.
- Microkernel** A modular kernel.
- Mirroring** Complete duplication of data to another disk; used by some levels of RAID.
- Modem** Modulator/Demodulator; modulates binary data into analog sound that can be carried on phone networks, and demodulates analog sound into binary data.
- Monoalphabetic cipher** A substitution cipher using one alphabet.

- Monolithic kernel** A statically compiled kernel.
- MOR** Minimum Operating Requirements; the minimum environmental and connectivity requirements for operating computer equipment.
- MTBF** Mean Time between Failures; the length of time a new or repaired system will run, on average, before failing.
- MTD** Maximum Tolerable Downtime; the total time a system can be inoperable before an organization is severely impacted.
- MTTR** Mean Time to Repair; the time required to recover a failed system.
- Multiprocessing** Multiple processes running on multiple CPUs.
- Multitasking** Multiple tasks (heavy-weight processes) running simultaneously on one CPU.
- NDA** Nondisclosure agreement; a contractual agreement ensuring that an individual or organization appreciates the legal responsibility to maintain confidentiality of sensitive information.
- Need to Know** A requirement that subjects have a specific need to know information for permission to access it.
- Nonrepudiation** Assurance that a specific user performed a specific transaction and that the transaction did not change.
- Object** A data file.
- CTAVE** Operationally Critical Threat, Asset, and Vulnerability Evaluation; a risk management framework developed by Carnegie Mellon University.
- OEP** Occupant Emergency Plan; a facility-based plan focused on safety and evacuation.
- OFDM** Orthogonal Frequency-Division Multiplexing; a newer wireless multiplexing method that allows simultaneous transmission using multiple independent wireless frequencies that do not interfere with each other.
- Offshoring** Outsourcing of services to an organization in another country.
- OLE** Object Linking and Embedding; a part of DCOM that links documents to documents.
- OOP** Object-Oriented Programming; a newer programming methodology that, unlike the older procedural programming methodology, treats a program as a series of connected objects that communicate via messages.
- Open system** A system of open hardware and standards that uses standard components from a variety of vendors.
- ORB** Object Request Broker; strategic middleware used to locate objects and communicate with them.
- OUI** Organizationally Unique Identifier; the first 24 bits of a MAC address.
- Outsourcing** The use of a third party to provide Information Technology support services previously performed in-house.
- PAN** Personal Area Network; a very small network with a range of 100 meters or less.
- PAP** Password Authentication Protocol; an insecure network authentication protocol that exposes passwords in cleartext.
- Parity** A means to achieve data redundancy without incurring the cost of mirroring in terms of disk usage and write performance.
- Patent** Intellectual property protection that grants a monopoly on the right to use, make, or sell an invention for a period of time.
- PCI-DSS** Payment Card Industry Data Security Standard; a standard created by the Payment Card Industry Security Standards Council (PCI SSC).
- PEAP** Protected EAP; an authentication protocol, similar to EAP-TTLS, that does not require client-side certificates.
- PGP** Pretty Good Privacy; software that integrates asymmetric, symmetric, and hash cryptography.
- PII** Personally Identifiable Information; data associated with a specific person, such as credit card information.
- Pipelining** A CPU feature that combines multiple steps into one combined process, allowing simultaneous fetch, decode, execute, and write steps for different instructions.
- PKI** Public Key Infrastructure; a system that leverages symmetric, asymmetric, and hash-based cryptography to manage digital certificates.
- Plaintext** An unencrypted message.
- Policy** High-level management directives; administrative control.

- Polyalphabetic cipher** A substitution cipher using multiple alphabets.
- Polyinstantiation** Allows two different objects with different data to have the same name.
- POP** Post Office Protocol; an email client protocol.
- PPP** Point-to-Point Protocol; a Layer 2 protocol that has largely replaced SLIP, adding confidentiality, integrity, and authentication.
- Private key** One half of an asymmetric key pair which must be kept secure.
- Procedure** A step-by-step guide for accomplishing a task; administrative control.
- Process** An executable program and its associated data loaded and running in memory.
- PROM** Programmable Read Only Memory; memory that can be written to once, typically at the factory.
- Public key** One-half of an asymmetric key pair which may be publicly posted.
- PVC** Permanent Virtual Circuit; a circuit that is always connected.
- Qualitative Risk Analysis** A risk analysis method that uses approximate values.
- Quantitative Risk Analysis** A risk analysis method that uses hard metrics such as dollars.
- RADIUS** Remote Authentication Dial In User Service; a UDP-based third-party authentication system.
- RAID** Redundant Array of Inexpensive Disks; the use of multiple disk drives to achieve greater data reliability, greater speed, or both.
- RAM** Random Access Memory; volatile hardware memory that loses integrity after loss of power.
- RBAC** Role-Based Access Control; the grouping of subjects into defined roles; access permissions are based on these roles, not on individual subjects.
- RC4** Rivest Cipher 4; a stream cipher used by WPA to provide confidentiality.
- Real evidence** Evidence consisting of tangible or physical objects.
- Return on Investment (ROI)** The amount of money saved by deploying a safeguard.
- RFID** Radio-Frequency Identification; a type of contactless card technology.
- Ring model** A form of CPU hardware layering that separates and protects domains (e.g., kernel mode and user mode) from each other.
- RISC** Reduced Instruction Set Computer; short, simple CPU instructions.
- Risk Analysis Matrix** A quadrant used to map the likelihood of a risk occurring against the consequences (or impact) that the risk would have.
- ROM** Read Only Memory; nonvolatile memory that maintains integrity after loss of power.
- Rootkit** Malware that replaces portions of the kernel and/or operating system.
- RPO** Recovery Point Objective; the amount of data loss or system inaccessibility (measured in time) that an organization can withstand.
- RSN** Robust Security Network; a part of 802.11i that allows changes to cryptographic ciphers as new vulnerabilities are discovered.
- RTO** Recovery Time Objective; the maximum time allowed to recover business or IT systems.
- S/MIME** Secure/Multipurpose Internet Mail Extension; an encryption standard that leverages PKI to encrypt and authenticate MIME-encoded email.
- SAML** Security Assertion Markup Language; an XML-based framework for exchanging security information, including authentication data.
- Sanction** An action taken as a result of a policy violation.
- SDLC (Applications)** Systems Development Life Cycle; a system development model that focuses on security in every phase.
- Security domain** The list of objects a subject is allowed to access.
- Servicemark** Intellectual property protection that allows the creation of a brand that distinguishes the source of services.
- SHA-1** A hash function that creates a 160-bit message digest.
- Shareware** Fully functional proprietary software that may be initially used free of charge. After a specific period of time, the shareware license typically requires payment.
- SLA** Service Level Agreement; a contractual agreement that helps ensure availability.
- SLE** Single Loss Expectancy; the cost of a single loss.
- SLIP** Serial Line Internet Protocol; a Layer 2 protocol that provides IP connectivity via asynchronous connections such as serial lines and modems.

- Smartcard** A physical access control device containing an integrated circuit. Also known as an *Integrated Circuit Card (ICC)*.
- SMDS** Switched Multimegabit Data Service; an older WAN technology that is similar to ATM.
- SMTP** Simple Mail Transfer Protocol; a store-and-forward protocol for exchanging email between servers
- SOCKS** A popular circuit-level proxy for routing network packets between clients and servers.
- Source code** Computer programming language instructions that are written in text that must be translated into machine code before execution by the CPU.
- Spring-bolt lock** A physical locking mechanism that “springs” in and out of a door jamb.
- SQL** Structured Query Language; the most popular database query language.
- SRAM** Static Random Access Memory; expensive and fast memory that uses small latches called “flip-flops” to store bits.
- SSH** Secure Shell; a secure replacement for Telnet, FTP, and the UNIX “R” commands.
- SSL** Secure Sockets Layer; a protocol for authenticating and providing confidentiality to network (e.g., web) traffic.
- Standard** The specific use of technology, often applied to hardware and software; administrative control.
- STP** Shielded Twisted Pair; network cabling that contains additional metallic shielding around each twisted pair of wires.
- Striping** A way to achieve performance gains by spreading data writes across multiple disks; used by some levels of RAID.
- Subject** An active entity on an Information System which accesses or changes data.
- SVC** Switched Virtual Circuit; a circuit that is established on demand.
- Swapping** The use of virtual memory to copy contents in primary memory (RAM) to or from secondary memory.
- Symmetric Encryption** Encryption that uses one key to encrypt and decrypt.
- T1** A dedicated 1.544-megabit circuit that carries 24 64-bit DS0 channels.
- T3** The equivalent of 28 Bundled T1s.
- TACACS** Terminal Access Controller Access Control System; a centralized access control system that requires users to send an ID and static (reusable) password for authentication.
- Tailgating** A method of entry in which an unauthorized person follows an authorized person into a building without providing credentials. Also known as *piggybacking*.
- TCO** Total Cost of Ownership, the cost of a safeguard.
- TCP** Transmission Control Protocol; a protocol that uses a three-way handshake to create reliable connections across a network.
- TCSEC** Trusted Computer System Evaluation Criteria, also known as the *Orange Book*; an evaluation model developed by the U.S. Department of Defense.
- Telnet** A network protocol that provides terminal emulation over a network.
- TFTP** Trivial File Transfer Protocol; a simple way to transfer files with no authentication or directory structure.
- Thin client** A simple computer system that relies on centralized applications and data.
- Thread** A light-weight process (LWP).
- Threat** A potentially negative occurrence.
- TKIP** Temporal Key Integrity Protocol; a protocol used by WPA to provide integrity.
- TLS** Transport Layer Security; the successor to SSL.
- TOCTOU** Time of Check/Time of Use; the altering of a condition after it has been checked by the operating system but before it is used.
- Trade secret** Business-proprietary information that is important to an organization’s ability to compete.
- Trademark** Intellectual property protection that allows for the creation of a brand that distinguishes the source of a product or products.
- Trojan** Malware that performs two functions: one benign (e.g., a game) and one malicious. Also called *Trojan Horse*.
- Tuple** A row in a relational database table.

UDP User Datagram Protocol; a simpler and faster cousin of TCP.

UTP Unshielded Twisted Pair; network cabling that uses pairs of wire twisted together.

Virtual memory Memory that provides virtual address mapping between applications and hardware.

Virus Malware that requires a carrier.

VPN Virtual Private Network; a method for sending private data over an insecure network such as the Internet.

Vulnerability A weakness in a system.

WAN Wide Area Network; a network typically covering cities, states, or countries.

Warded lock A type of lock that requires the turn of a key through channels (called wards) to be opened.

WEP Wired Equivalent Privacy; a very weak 802.11 security protocol.

WORM Write Once Read Many; memory that can be written to once and read many times.

Worm Malware that self-propagates.

WPA Wi-Fi Protected Access; a partial implementation of 802.11i.

WPA2 Wi-Fi Protected Access 2; the full implementation of 802.11i.

WRT Work Recovery Time; the time required to configure a recovered system.

X.25 An older packet-switched WAN protocol.

XML eXtensible Markup Language; a markup language designed as a standard for encoding documents and data.

Index

Note: Page number followed by *b* indicate boxes; *f* indicate figures; *t* indicate tables.

A

- 802.11
 - definition, 123
 - types, 123*t*
- 802.11i, 124
- 802.1X, 121
- AAA, *see* Accountability;
Authentication; Authorization
- Abstraction, 70
- Access control
 - assessment, 33–34
 - authentication methods, 26–31
 - categories and types, 25–26
 - centralized, 22
 - CIA triad, 19–20, 20*f*
 - cornerstone concepts, 19
 - DAC, 22
 - decentralized, 22
 - defense in depth, 21
 - definition, 19–21
 - identification and AAA, 20–21
 - least privilege, 21
 - MAC, 22
 - models, 21–24
 - need to know, 21
 - nonrepudiation, 21
 - procedural issues, 24–25
 - protocols and frameworks, 23–24
 - subjects and objects, 21
 - technologies, 31–33
- Access Control Lists, *see* ACL
- Access control matrix
 - definition, 82–83
 - permissions, 83*t*
- Accountability
 - and Diameter, 23
 - definition, 21
- Accreditation
 - basic considerations, 83–85
 - definition, 15–16
- ACL (Access Control Lists), 25
- “Act honorably, honestly, justly, responsibly, and legally,” 171
- Adaptive chosen ciphertext, 47
- Adaptive chosen plaintext, 47
- Administrative controls, 25
- Administrative law, 163
- Administrative security
 - background checks, 148
 - change management, 148–149
 - configuration management, 148
 - definition, 147–149
 - NDA, 148
 - rotation of duties/job rotation, 148
 - separation of duties, 147
- ADSL (Asymmetric Digital Subscriber Line)
 - definition, 125
 - speed and distance, 125*t*
- “Advance and protect the profession,” 172
- AES (Advanced Encryption Standard)
 - cryptographic strength, 40
 - definition, 44
 - finalists, 44*t*
 - WPA2, 124
- Aggregation, 80
- AH (Authentication Header)
 - definition, 50–51
 - IPsec, 122
 - transport mode, 51
- Alarms, as perimeter defense, 58
- ALE (Annualized Loss Expectancy)
 - calculation, 4–5
 - definition, 5
 - encrypted laptops, 6*t*
 - Quantitative Risk Analysis as, 8
 - risk equations, 5*t*
 - unencrypted laptops, 6*t*
- Allocated space, 166
- ALU (Arithmetic Logic Unit), 72
- AMP (Asymmetric Multiprocessing), 73
- Annual Rate of Occurrence, *see* ARO
- Annualized Loss Expectancy, *see* ALE
- Applet, 79
- Application development methods
 - OOP, 134–136
 - overview, 131–133
 - RAD, 133
 - SDLC, 133
 - Spiral Model, 131, 132*f*
 - Waterfall Model, 131
 - XP, 133
- Application development security
 - CMM, 138–139
 - databases, 139–143
 - disclosure, 138
 - overview, 129
 - programming concepts, 129–130
 - software vulnerabilities, 137–138
- Application layer
 - definition, 112, 113
 - proxy firewalls, 119, 120
 - TCP/IP protocols, 115–116
- Arithmetic Logic Unit, *see* ALU
- ARO (Annual Rate of Occurrence)
 - definition, 4
 - encrypted laptop ALE, 6*t*
 - risk equations, 5*t*
 - unencrypted laptop ALE, 6*t*
- AS (Authentication Server), for SSO, 31
- Assembly language, 129–130
- Assessment step, disaster recovery process, 93
- Asset, 2
- Asset Value, *see* AV (Asset Value)
- Asymmetric Digital Subscriber Line, *see* ADSL
- Asymmetric Encryption
 - basic concept, 45–46
 - definition, 41
 - methods, 45–46
 - vs.* symmetric, 46
- Asymmetric Multiprocessing, *see* AMP
- Asynchronous Dynamic Token, as Type 2 authentication, 28
- ATM (Asynchronous Transfer Mode), 110, 117
- Attribute
 - and semantic integrity, 141
 - definition, 139–140
- Attribute Value Pairs, *see* AVP
- Auditing, 13–15

- Authentication
 - and Diameter, 23
 - and identification, 20
 - in cryptography, 40
 - methods, 26–31
 - Type 1, 26–27
 - Type 2, 27–28
 - Type 3, 28–30
 - Authentication Header, *see* AH
 - Authentication protocol, 120–121
 - Authentication Server, *see* AS
 - Authorization
 - and Diameter, 23
 - definition, 20
 - AV (Asset Value)
 - definition, 4
 - encrypted laptop ALE, 6t
 - risk equations, 5t
 - unencrypted laptop ALE, 6t
 - Availability
 - in CIA triad, 19–20, 20f
 - definition, 20
 - AVP (Attribute Value Pairs), and Diameter, 23
- B**
- Backdoor, 78
 - Background checks, 148
 - Backup, as fault tolerance, 151
 - “Bad” blocks/clusters/sectors, 166
 - Baseband, 110
 - Baseline
 - definition, 11
 - example, 11t
 - Basic Input Output System, *see* BIOS
 - BCI (Business Continuity Institute), BCP/DRP frameworks, 104
 - BCM (Business Continuity Management), 104
 - BCP (Business Continuity Plan)
 - basic considerations, 89–93
 - BCI framework, 104
 - BIA, 94–96
 - BS-25999 framework, 104
 - Call Tree, 99
 - change management, 102
 - critical state assessment, 94
 - definition, 89
 - development, 93–99
 - disaster recovery process, 92–93
 - disasters, 90
 - DRP relationship, 90
 - frameworks, 103–104
 - goals and focus, 90
 - ISO/IEC-27031 framework, 103–104
 - maintenance, 102
 - mistakes, 102
 - NIST SP 800-34 framework, 103
 - plans summary, 99t
 - preventive control identification, 97
 - project initiation, 94
 - recovery strategy, 97–98
 - related plans, 91f, 98
 - Bell-LaPadula model
 - and Biba model, 82b
 - definition, 80
 - rules and properties, 81b
 - BIA (Business Impact Analysis)
 - BCP/DRP-focused risk assessment, 95
 - critical asset identification, 95
 - definition, 94–96
 - failure/recovery metrics, 95–96
 - MOR, 96
 - MTBF, 96
 - MTD determination, 95
 - MTTR, 96
 - risk mitigation, 97
 - RPO, 96
 - RTO/WRT, 96
 - Biba model
 - and Bell-LaPadula model, 82b
 - definition, 81–82
 - rules, 82b
 - Biometric systems
 - accuracy, 28–29
 - control types, 29–30
 - enrollment and throughput, 28
 - BIOS (Basic Input Output System), 76
 - Birthday attack, 48
 - Bits, 111
 - Black-box test, *see* Zero knowledge test
 - Block ciphers
 - AES, 44
 - definition, 41
 - Blowfish, 44
 - Bollard, as perimeter defense, 56
 - Boot sector virus, 78
 - Bridges
 - definition, 117–118
 - example, 118f
 - Broadband, 110
 - BRP (Business Recovery Plan)
 - as BCP plan, 98
 - definition, 99t
 - Brute force attack, 47
 - BS-25999, BCP/DRP frameworks, 104
 - Buffer Overflow, 137
 - Buildings, site configuration issues, 60
 - Business Continuity Institute, *see* BCI
 - Business Continuity Management, *see* BCM
 - Business Continuity Plan, *see* BCP
 - Business Impact Analysis, *see* BIA
 - Business interruption, DRP testing, 101
 - Business Owner, *see* Data Owner
 - Business Recovery Plan, *see* BRP
 - Bytecode, 130
- C**
- Cable modems, 125
 - Cache memory, 74
 - Call Tree
 - definition, 99
 - training/testing, 101
 - Capability Maturity Model, *see* CMM
 - Carbon dioxide, as fire suppression agent, 65
 - CAs (Certificate Authorities), 50
 - Cathode Ray Tube cameras, *see* CRT cameras
 - CBC (Cipher Block Chaining), 42, 43t
 - CBC-MAC (Cipher Block Chaining Message Authentication Code), 50
 - CCD (Charge Couple Discharge), and CCTV, 56
 - CCMP (Counter Mode CBC MAC Protocol), WPA2, 124
 - CCTV (Closed-Circuit Television)
 - detective controls, 25
 - as perimeter defense, 56
 - Ceilings, as perimeter defense, 59
 - Central Processing Unit, *see* CPU
 - Centralized access control, 22
 - CER (Crossover Error Rate)
 - definition, 29
 - example, 29f
 - Certificate Authorities, *see* CAs
 - Certificate Revocation Lists, *see* CRLs
 - Certification
 - basic considerations, 83–85
 - definition, 15–16
 - CFB (Cipher Feedback), 43, 43t
 - Chain of Custody form, 168f

- Change management
 - administrative security, 148–149
 - BCP/DRP maintenance, 102
 - general flow, 149
- CHAP (Challenge Handshake Authentication Protocol), 23–24, 120–121
- Charge Couple Discharge, *see* CCD
- Checklist testing, DRP, 100
- Chinese Wall model, 82
- Chosen ciphertext, 47
- Chosen plaintext, 47
- CIA (Confidentiality, Integrity, and Availability) triad
 - and DAD, 20
 - definition, 19–20
 - pyramid example, 20*f*
- Cipher, 39
- Cipher Block Chaining, *see* CBC
- Cipher Block Chaining Message Authentication Code, *see* CBC-MAC
- Cipher Feedback, *see* CFB
- Ciphertext
 - and strong encryption, 40
 - cryptographic attacks, 47
 - definition, 39
- Circuit-level proxies, 120
- Circuit-switched networks, 110
- Circumstantial evidence, 167
- CISC (Complex Instruction Set Computer), 73
- Civil law
 - definition, 162–163
 - as legal system, 161
- Clark-Wilson model, 82
- Classes, and OOP, 134, 135
- Classes of fires, 63–64, 64*t*
- Clearance, Access Control
 - procedures, 24
- Clipper Chip, 52
- Clipping level, 156
- Closed-Circuit Television, *see* CCTV
- Closed system, 71
- Closed-source software, 130
- CMM (Capability Maturity Model)
 - definition, 138–139
 - levels, 139
- CMP (Crisis Management Plan), as BCP plan, 98
- Coaxial cable, and EMI, 62
- COBIT (Control Objectives for Information and related Technology), 14–15
- Cold site, 98
- Collusion, 147
- COM (Component Object Model), 136
- Combination locks, 57
- Common Criteria
 - definition, 84–85
 - levels, 85
- Common law
 - administrative branch, 163
 - civil branch, 162–163
 - criminal branch, 162
 - as legal system, 162
- Common Object Request Broker Architecture, *see* CORBA
- Communication
 - network concepts, 109
 - secure, 120–125
- Communication step, disaster
 - recovery process, 93
- Compensating control, 26
- Compensatory damages, 163*t*
- Compiler, 130
- Complex Instruction Set Computer, *see* CISC
- Compliance, security policy
 - definition, 10
- Component Object Model, *see* COM
- Computer attacks
 - malware, 156
 - password guessing/cracking, 156
- Computer bus
 - definition, 71
 - simplified example, 71*f*
- Computer crime
 - legal concepts, 164
 - perpetration, 164
- Computer Ethics Institute, 170
- Computer Fraud and Abuse Act, 167
- Conficker worm, as threat example, 2
- Confidentiality
 - and privacy, 12–13
 - in CIA triad, 19–20, 20*f*
 - in cryptography, 40
 - definition, 19
 - WPA2, 124
- Configuration management, 148
- Constrained user interface, 142
- Continuity of Operations
 - fault tolerance, 151–155
 - SLA, 150
- Continuity of Operations Plan, *see* COOP
- Continuity of Support Plan
 - as BCP plan, 98
 - definition, 99*t*
- Contraband checks, as perimeter
 - defense, 58
- Control frameworks, 13–15
- Control Objectives for Information and related Technology, *see* COBIT
- Control unit, 72
- Controls, *see* Safeguards
- COOP (Continuity of Operations Plan)
 - and BCP–DRP relationship, 90
 - as BCP plan, 98
 - definition, 99*t*
- Copyright, legal issues, 165
- CORBA (Common Object Request Broker Architecture), 136
- Corrective controls, 26
- Corrosion, environmental controls, 63
- Council of Europe Convention on Cybercrime, 164
- Counter, *see* CTR (Counter)
- Counter Mode CBC MAC Protocol, *see* CCMP
- Covert channel, 77–78
- CPU (Central Processing Unit) and components, 71–73
- Crime
 - computer, legal concepts, 164
 - site selection issues, 60
- Criminal law, 162
- Crippleware, 130
- Crisis Communications Plan
 - Call Tree, 99
 - definition, 99*t*
- Crisis Management Plan, *see* CMP
- Critical assets, identification, 95
- Critical state, assessment, 94
- CRLs (Certificate Revocation Lists), 50
- Crossover Error Rate, *see* CER
- Cross-Site Request Forgery, *see* CSRF
- Cross-Site Scripting, *see* XSS
- Crosstalk, mitigation, 62
- CRT (Cathode Ray Tube) cameras, 57*b*
- Cryptanalysis
 - cryptographic attacks, 48
 - definition, 39
- Cryptography, *see also* Encryption
 - attack definitions and types, 46–48
 - cornerstone concepts, 39–41
 - definition, 39
 - strength, 40

- Cryptography implementation
 - basic considerations, 48–52
 - CBC-MAC, 50
 - digital signatures, 48–49
 - escrowed encryption, 52
 - HMAC, 49
 - IPsec, 50–51
 - PGP, 52
 - PKI, 50
 - S/MIME, 52
 - SSL/TLS, 51–52
- Cryptology, 39
- Crystal-box test, *see* Full knowledge test
- C/S session key, Kerberos steps, 32
- CSRF (Cross-Site Request Forgery), 137
- CTR (Counter), 43, 43*t*
- Custodian, role and responsibility, 12
- Customary law, as legal system, 162
- Cyber Incident Response Plan
 - as BCP plan, 98
 - definition, 99*t*
- D**
- DAC (Discretionary Access Control), 22
- DAD (Disclosure, Alteration, and Destruction), and CIA triad, 20
- Data Definition Language, *see* DDL
- Data Encryption Algorithm, *see* DEA
- Data Encryption Standard, *see* DES
- Data integrity, 20
- Data link layer, 111
- Data Manipulation Language, *see* DML
- Data mining, 80
- Data Owner, role and responsibility, 12
- Data remanence, media security, 149
- Database Administrators, *see* DBAs
- Database Management System, *see* DBMS
- Database security, basic considerations, 80
- Databases
 - example table, 140*t*
 - integrity, 141*t*, 142
 - normalization, 141
 - query languages, 142
 - relational, 139–142, 140*t*
 - replication and shadowing, 142–143
 - views, 142, 142*t*
- DBAs (Database Administrators), 139
- DBMS (Database Management System)
 - data table integrity, 140–141
 - definition, 139
- DCOM (Distributed Component Object Model), 136
- DDL (Data Definition Language), 142
- DDoS (Distributed Denial of Service), 156
- DEA (Data Encryption Algorithm), 41
- Decentralized access control, 22
- Decryption, 39
- Defense-in-depth, 21
- Degaussing, media security, 150
- Delegation, and OOP, 134, 135*f*
- Deluge systems, for fire suppression, 66
- Demarc, 60
- Denial of Service, *see* DoS
- DES (Data Encryption Standard)
 - algorithm definition, 42
 - definition, 41–44
 - modes, 42–43
 - single DES, 43
- Destruction of data, media security, 149–150
- Detective controls, 25
- Deterrent controls, 26
- Diameter, 23
- Differential backup, 151
- Differential cryptanalysis, 48
- Diffie–Hellman Key Agreement Protocol, method, 45
- Digital forensics, legal issues, 166
- Digital signatures
 - creation, 49*f*
 - definition, 48–49
 - verification, 49*f*
- Digital Subscriber Line, *see* DSL
- Digital Video Recorder, *see* DVR
- Direct evidence, 167
- Direct Sequence Spread Spectrum, *see* DSSS
- Directive controls, *see* Administrative controls
- Directory Path Traversal, 137
- Disaster Recovery Planning, *see* DRP
- Disaster recovery process
 - assessment, 93
 - communication, 93
 - reconstitution, 93
 - response, 92
 - team activation, 92–93
- Disasters
 - categories and causes, 91
 - overview, 90
- Disclosure, Alteration, and Destruction, *see* DAD
- Disclosure, software vulnerabilities, 138
- Discrete logarithm, method, 45
- Discretionary Access Control, *see* DAC
- Diskless workstation, 77
- Disruptive events
 - examples, 92*t*
 - overview, 90
 - types, 92
- Distributed access control, *see* Decentralized access control
- Distributed Component Object Model, *see* DCOM
- Distributed Denial of Service, *see* DDoS
- DML (Data Manipulation Language), 142
- DNS (Domain Name System)
 - definition, 115
 - reflection attack, 157*t*
- Documentation
 - basic considerations, 9–11
 - examples, 11*t*
- Dogs, as perimeter defense, 59
- Domain Name System, *see* DNS
- Doors, as perimeter defense, 58
- DoS (Denial of Service)
 - as availability attack, 20
 - definition, 156
 - examples, 157*t*
- DRAM (Dynamic Random Access Memory), 74
- Drive encryption, system defenses, 61
- DRP (Disaster Recovery Planning)
 - basic considerations, 89–93
 - BCI framework, 104
 - BCP relationship, 90
 - BIA, 94–96
 - BS-25999 framework, 104
 - Call Tree, 99
 - change management, 102
 - critical state assessment, 94
 - definition, 89, 99*t*
 - development, 93–99
 - disaster recovery process, 92–93
 - disasters, 90
 - frameworks, 103–104
 - goals and focus, 90
 - ISO/IEC-27031 framework, 103–104

- maintenance, 102
 - mistakes, 102
 - NIST SP 800-34 framework, 103
 - preventive control identification, 97
 - project initiation, 94
 - recovery strategy, 97–98
 - testing, 100–101
 - training, 101
 - Dry pipe systems, for fire suppression, 66
 - Dry powder, as fire suppression agent, 64
 - DSL (Digital Subscriber Line)
 - definition, 124–125
 - speed and distances, 125*t*
 - DSSS (Direct Sequence Spread Spectrum), 122
 - Dual stack systems, 114
 - Due care, 164
 - Due diligence, 164
 - Dumpster diving, 61
 - DVR (Digital Video Recorder), and CCTV, 56
 - Dynamic passwords, as Type 1 authentication, 27
 - Dynamic Random Access Memory, *see* DRAM
 - Dynamic signature, biometric controls, 30
- E**
- E1, 116
 - E3, 116
 - EAL (Evaluation Assurance Level), 84, 85
 - EAP (Extensible Authentication Protocol)
 - definition, 121
 - types, 121
 - EAP-TLS (EAP–Transport Layer Security), 121
 - EAP-TTLS (EAP–Tunneled Transport Layer Security), 121
 - ECB (Electronic Code Book), 42, 43*t*
 - ECC (Elliptic Curve Cryptography), method, 45–46
 - EDE (Encrypt, Decrypt, Encrypt), TDES, 44
 - EEPROM (Electrically-Erasable Programmable Read Only Memory), 75
 - EER (Equal Error Rate), 29
 - EES (Escrowed Encryption Standard), Clipper Chip, 52
 - EF (Exposure Factor)
 - definition, 4
 - encrypted laptop ALE, 6*t*
 - risk equations, 5*t*
 - unencrypted laptop ALE, 6*t*
 - Electrically-Erasable Programmable Read Only Memory, *see* EEPROM
 - Electricity
 - common faults, 62
 - environmental controls, 61–63
 - Electromagnetic Interference, *see* EMI
 - Electronic Code Book, *see* ECB
 - Elliptic Curve Cryptography, *see* ECC
 - Emergency power, DRP training, 101
 - EMI (Electromagnetic Interference), environmental controls, 62
 - Encapsulating Security Payload, *see* ESP
 - Encrypt, Decrypt, Encrypt, *see* EDE
 - Encryption, *see also* Cryptography
 - definition, 39
 - drives/tapes, 61
 - pattern destruction, 40
 - End-User License Agreements, *see* EULAs
 - Enrollment, 28
 - Entity integrity
 - and tuple, 141
 - definition, 140–141
 - Environmental controls
 - definition, 61–64
 - electricity, 61–63
 - evacuation notes, 63
 - heat/flame/smoke detectors, 63
 - Environmental threats, 91
 - EPROM (Erasable Programmable Read Only Memory), 75
 - Equal Error Rate, *see* EER
 - Escrowed Encryption Standard, *see* EES
 - ESP (Encapsulating Security Payload)
 - definition, 51
 - IPsec, 122
 - tunnel mode, 51
 - Ethernet
 - basic technology, 116
 - Layer 2, 111
 - Ethics
 - Computer Ethics Institute, 170
 - definition, 170–172
 - IAB, 171
 - (ISC)² © Code of Ethics, 171–172
 - EULAs (End-User License Agreements), legal issues, 165
 - European Union Data Protection Directive
 - overview, 169–170
 - principles, 169
 - European Union Privacy, 169
 - European Union–United States Safe Harbor, 169–170
 - Evacuation notes, environmental controls, 63
 - Evaluation Assurance Level, *see* EAL
 - Evaluation methods
 - basic considerations, 83–85
 - International Common Criteria, 84–85
 - ITSEC, 84
 - Orange Book, 83
 - PCI-DSS, 85
 - Evidence
 - integrity, 167
 - legal issues, 166
 - Exam tips
 - BIA, 95*b*
 - CRT cameras, 57*b*
 - DES, 42*b*
 - (ISC)² © Code of Ethics, 170*b*
 - multitasking/multiprogramming/multithreading, 73*b*
 - OSI vs. TCP/IP, 113
 - Quantitative Risk Analysis, 7*b*
 - RAID, 153
 - retina scans, 30
 - risk formula, 3*b*
 - virtual memory, 75*b*
 - Waterfall Model, 131
 - Exclusive Or, *see* XOR
 - Exposure Factor, *see* EF
 - Extensible Authentication Protocol, *see* EAP
 - Extensible Authentication Protocol–Transport Layer Security, *see* EAP-TLS
 - Extensible Authentication Protocol–Tunneled Transport Layer Security, *see* EAP-TTLS
 - eXensible Markup Language, *see* XML
 - Extranet, 110
 - eXtreme Programming, *see* XP

F

Facial scan, biometric controls, 30
 Failure metrics, BIA, 95–96
 False Accept Rate, *see* FAR
 False Reject Rate, *see* FRR
 FAR (False Accept Rate), 28
 Fault tolerance
 backup, 151
 definition, 151–155
 RAID, 151–154
 system redundancy, 154–155
 FDX (Fetch, Decode, Execute), 72
 Fences, as perimeter defense, 55
 FHSS (Frequency Hopping Spread Spectrum), 122
 Fiber optic cable, and EMI, 62
 File Transfer Protocol, *see* FTP
 Financial damages, types, 163*t*
 Fingerprint minutiae
 definition, 29
 example, 29*f*
 Fingerprints, biometric controls, 29
 Fire classess, 63–64, 64*t*
 Fire extinguishers, for fire
 suppression, 66
 Fire rating, walls, 59
 Fire suppression agents, 64–66, 64*t*
 Firewalls, 118–120
 Firmware, 75–76
 Flame detectors, environmental
 controls, 63
 Flash memory, 76
 Floors, as perimeter defense, 59
 Foreign keys
 definition, 140
 and referential integrity, 141
 “Forgotten password” procedure, 10
 Fraggle, 157*t*
 Frame relay, 117
 Frameworks
 access control, 23–24
 BCP/DRP, 103–104
 secure communications, 120–121
 Free software, 130
 Freeware, 130
 Frequency Hopping Spread
 Spectrum, *see* FHSS
 FRR (False Reject Rate), 28
 FTP (File Transfer Protocol)
 application layer, 113
 application layer proxies, 120
 definition, 115
 Full backup, 151
 Full Disclosure, software
 vulnerabilities, 138

Full knowledge test, 34

Full-duplex communication, 109

G

GAN (Global Area Network), 110
 Gates
 classes, 56
 as perimeter defense, 55
 Generators, environmental controls,
 62
 GLB (Greatest Lower Bound), 80–81
 GLBA (Gramm-Leach-Bliley Act),
 and offshoring, 13
 Global Area Network, *see* GAN
 Global Positioning System, *see* GPS
 Good Practice Guidelines, *see* GPG
 GPG (Good Practice Guidelines),
 104
 GPS (Global Positioning System),
 biometric controls, 31
 Gramm-Leach-Bliley Act, *see* GLBA
 Greatest Lower Bound, *see* GLB
 Greenfield, 59
 Guards, as perimeter defense, 59
 Guideline
 definition, 11
 example, 11*t*

H

HA (High Availability) clusters
 database replication and
 shadowing, 142–143
 definition, 155
 Half-duplex communication, 109
 Halon, as fire suppression agent, 65
 Halon substitutes, as fire suppression
 agent, 65, 65*b*
 Hand geometry, biometric controls,
 30
 Hard-coded credentials, 137
 Hardware redundancy, 155
 Hardware segmentation, 75
 Hash Function, 46
 Hashed Message Authentication
 Code, *see* HMAC
 Hashing, 41
 HAVAL (Hash of Variable Length),
 46
 HDSL (High Data Rate Digital
 Subscriber Line)
 definition, 125
 speed and distance, 125*t*
 Health Insurance Portability and
 Accountability Act, *see*
 HIPAA

Hearsay, 167

Heat detectors, environmental
 controls, 63

Heating, Ventilation, and Air
 Conditioning, *see* HVAC

Heavy-Weight Process, *see* HWP

High Availability clusters, *see* HA
 clusters

High Data Rate Digital Subscriber
 Line, *see* HDSL

HIPAA (Health Insurance Portability
 and Accountability Act), and
 offshoring, 13

HMAC (Hashed Message
 Authentication Code), 49

Host-to-host transport layer, 113

Hot site, 97

HTML (HyperText Markup
 Language)
 and CSRF, 138
 definition, 116
 vs. XML, 79

HTTP (HyperText Transfer Protocol)
 application layer proxies, 120
 definition, 116
 and SSH, 115
 SSL, 122

HTTPS (HyperText Transfer Protocol
 Secure)
 definition, 116
 SSL/TLS, 51, 122

Hubs, 117

Human disasters, 91

Humidity, environmental
 controls, 63

HVAC (Heating, Ventilation, and Air
 Conditioning), environmental
 controls, 62

HWP (Heavy-Weight Process), 73

HyperText Markup Language, *see*
 HTML

HyperText Transfer Protocol, *see*
 HTTP

HyperText Transfer Protocol Secure,
see HTTPS

I

IAB (Internet Activities Board),
 ethics, 171

ICC (Integrated Circuit Card), *see*
 Smart card

ICMP (Internet Control Message
 Protocol)

function, 114–115

Layer 3, 111

- Identification
 - and AAA, 20–21
 - and authentication, 20
 - critical assets, 95
 - preventive controls, 97
 - Identity Theft Enforcement and Restitution Act of 2008, 167
 - IDL (Interface Definition Language), and CORBA, 136
 - IKE (Internet Key Exchange), 51
 - IMAP (Internet Message Access Protocol), 115
 - Impact, 3
 - Incremental backup, 151
 - Inference, 80
 - Information Owner, *see* Data Owner
 - Information Security Governance
 - auditing and control frameworks, 13–15
 - certification and accreditation, 15–16
 - definition, 9–16
 - documentation, 9–16, 11*t*
 - outsourcing and offshoring, 13
 - overview, 1
 - policy, 9–16
 - privacy, 12–13
 - roles and responsibilities, 12
 - security awareness and training, 12
 - Information Technology
 - Infrastructure Library, *see* ITIL
 - Information Technology Security Evaluation Criteria, *see* ITSEC
 - Inheritance, and OOP, 134, 135
 - Integrated Services Digital Network, *see* ISDN
 - Integrity
 - in CIA triad, 19–20, 20*f*
 - in cryptography, 40
 - databases, 140–141, 141*t*, 142
 - definition, 20
 - evidence, 167
 - types, 20
 - Integrity Axiom, Biba model, 82
 - Integrity models, types, 81–82
 - Intellectual property, legal issues, 164–165
 - Interface Definition Language, *see* IDL
 - International Data Encryption Algorithm, 44
 - Internet
 - and ethics, 171
 - definition, 110
 - protocols, 113–115
 - Internet Activities Board, *see* IAB
 - Internet Control Message Protocol, *see* ICMP
 - Internet Key Exchange, *see* IKE
 - Internet layer, 112–113
 - Internet Message Access Protocol, *see* IMAP
 - Internet Protocol Security, *see* IPsec
 - Internet Protocol version 4, *see* IPv4
 - Internet Protocol version 6, *see* IPv6
 - Internet Security Association and Key Management Protocol, *see* ISAKMP
 - Interpreted code, 130
 - Interrupt, 72–73
 - Intranet, 110
 - Investigations, legal issues, 165–167
 - IPsec (Internet Protocol Security)
 - definition and protocols, 50–51
 - secure communications, 122
 - IPv4 (Internet Protocol version 4)
 - dual stack systems, 114
 - function, 113–114
 - IPsec, 122
 - Layer 3, 111
 - IPv6 (Internet Protocol version 6)
 - dual stack systems, 114
 - function, 114
 - IPsec, 122
 - Layer 3, 111
 - Iris scan, biometric controls, 29–30
 - ISAKMP (Internet Security Association and Key Management Protocol), 51
 - (ISC)² © Code of Ethics, 171–172
 - ISDN (Integrated Services Digital Network), 124
 - ISO 17799, *see* ISO 27002
 - ISO 27000 series, 14
 - ISO 27002
 - basic areas, 14
 - definition, 14
 - ISO/IEC-27031
 - BCP/DRP frameworks, 103–104
 - purpose, 103
 - IT Contingency Plan, 99*t*
 - ITIL (Information Technology Infrastructure Library), 15
 - ITSEC (Information Technology Security Evaluation Criteria), 84
- J**
- Job rotation, 148
- K**
- KDC (Key Distribution Center), Kerberos steps, 32
 - Kerberos
 - characteristics, 31
 - components, 31
 - definition, 31–33
 - operational steps, 32–33, 32*f*
 - Kernel, 76–77
 - Kernel-mode rootkit, 79
 - Key Distribution Center, *see* KDC
 - Key locks, 57
 - Keyboard dynamics, biometric controls, 30
 - Keying options, TDES, 44
 - Known key, 47
 - Known plaintext, 47
- L**
- Labels, Access Control procedures, 24
 - LAN (Local Area Network)
 - definition, 110
 - technologies and protocols, 116
 - Land attack, 157*t*
 - Laptops, ALE, 6*t*
 - Lattice-based Access Controls
 - definition, 80–81
 - example, 81*f*
 - Laws and Regulations
 - Computer Fraud and Abuse Act, 167
 - PATRIOT Act, 168
 - privacy, 169–170
 - Layer 1, 111
 - Layer 2
 - bridges, 117–118
 - definition, 111
 - Layer 3
 - application-layer proxy firewalls, 119
 - definition, 111
 - firewalls, 118–119
 - ICMP, 114–115
 - Layer 4
 - application-layer proxy firewalls, 119
 - definition, 111
 - firewalls, 118–119
 - Layer 5, 111
 - Layer 6, 112

Layer 7, 112
 Layered defenses, *see* Defense-in-depth
 Layering
 components, 70*b*
 definition, 69
 LEAP (Lightweight Extensible Authentication Protocol), 121
 Least privilege, 21
 Least Upper Bound, *see* LUB
 Legal issues
 computer crime, 164
 digital forensics, 166
 due care/diligence, 164
 evidence, 166, 167
 evidence integrity, 167
 intellectual property, 164–165
 international cooperation, 164
 investigations, 165–167
 Legal systems
 civil law, 161
 common law, 162
 major systems, 161–162
 religious and customary law, 162
 Lights, as perimeter defense, 56
 Lightweight Extensible Authentication Protocol, *see* LEAP
 Light-Weight Process, *see* LWP
 Linear cryptanalysis, 48
 Linux system, vulnerability example, 2
 Local Area Network, *see* LAN
 Locks, types, 57
 Logic bomb, 79
 LUB (Least Upper Bound), 80–81
 Lumen, 56
 Lux, 56
 LWP (Light-Weight Process), 73

M

MAAs (Mutual Aid Agreements), *see* Reciprocal agreement
 MAC (Access Control), 22
 MAC (Telecommunications)
 bridges, 117–118
 Layer 2, 111
 MAC address, 113
 Machine code, 129–130
 Macro virus, 78, 157*t*
 MAD (Maximum Allowable Downtime), *see* MTD
 Magnetic stripe cards, 57–58
 Maintenance hook, 78
 Malware (Malicious code)
 definition, 78–79, 156
 types, 157*t*
 MAN (Metropolitan Area Network), 110
 Man in the Middle, *see* MITM
 Mandatory Access Control, *see* MAC
 Mantrap, 58
 MAO (Maximum Acceptable Outage), *see* MTD
 Marking, site design issues, 60
 Mathematic operators, and OOP, 134
 Maximum Tolerable Downtime, *see* MTD
 MD5 (Message Digest 5), 46
 Mean Time Between Failure, *see* MTBF
 Mean Time to Repair, *see* MTTR
 Media
 cleaning/destruction, 61
 sanitization, 149–150
 security, 149–150
 storage/transportation, 61
 Meet-in-the-middle attack, 47
 Memory
 protection, 74–76
 types, 73–74
 Message Digest 5, *see* MD5
 Message Integrity Check, *see* MIC
 Messages, and OOP, 134, 135
 Methods, and OOP, 134, 135
 Metropolitan Area Network, *see* MAN
 MIC (Message Integrity Check), WPA2, 124
 Microkernel, 76
 Microsoft Windows system, vulnerability example, 2
 MIMO (Multiple-Input Multiple-Output), 802.11, 123
 Minimum Operating Requirements, *see* MOR
 Mirroring, 152
 MITM (Man in the Middle), 156
 Mobile site, 98
 Modem (Modulator/Demodulator), 120
 Modified Waterfall Model
 definition, 131
 example, 132*f*
 Modulator/Demodulator, *see* Modem
 Monkey in the Middle, *see* MITM
 Monoalphabetic cipher, 40

Monolithic kernel, 76
 Montreal Accord, 65
 MOR (Minimum Operating Requirements), 96
 Motherboard, 71
 Motion detectors, as perimeter defense, 58
 MTBF (Mean Time Between Failure), 96
 MTD (Maximum Tolerable Downtime)
 alternate terms, 95
 BIA goals, 94–95
 determination, 95
 MTO (Maximum Tolerable Outage), *see* MTD
 MTTR (Mean Time to Repair), 96
 Multi-factor authentication, *see* Strong authentication
 Multipartite virus, 79
 Multiple-Input Multiple-Output, *see* MIMO
 Multiprocessing, 73
 Multitasking
 definition, 73
 vs. multiprogramming/
 multithreading, 73*b*
 Multithreading, 73*b*

N

National Computer Security Center, *see* NCSC
 National Institute of Standards and Technology, *see* NIST
 National Security Agency, *see* NSA
 Natural disasters, 91
 NCSC (National Computer Security Center), Orange Book, 83
 NDA (Nondisclosure Agreement)
 administrative security, 148
 legal issues, 165
 Need to know, 21
 Network access layer
 definition, 112
 TCP/IP model, 113–115
 Network Architecture and Design
 application-layer TCP/IP protocols, 115–116
 basic concepts, 109–110
 LAN, 116
 network access, 113–115
 OSI model, 110–112, 111*t*
 OSI *vs.* TCP/IP, 112*t*
 TCP/IP model, 112–113
 WAN, 116–117

- Network attacks
 - DoS/DDoS, 156, 157*t*
 - malware, 156
 - MITM, 156
 - session hijacking, 156
 - Network devices
 - bridges, 117–118, 118*f*
 - firewalls, 118–120
 - modem, 120
 - overview, 117–120
 - repeaters and hubs, 117
 - routers, 118
 - switches, 118
 - Network Interface Cards, *see* NICs
 - Network layer, 111
 - Network protocol analyzer, 118
 - Network Video Recorder, *see* NVR
 - NICs (Network Interface Cards)
 - firewalls, 118–119
 - hardware redundancy, 155
 - network access layer, 112
 - NIST (National Institute of Standards and Technology)
 - Orange Book, 83
 - Special Publication 800-12, policy types, 10
 - Special Publication 800-14, application development, 133
 - Special Publication 800-34
 - BCP plans, 98, 99*t*
 - BCP/DRP development, 93
 - BCP/DRP frameworks, 103
 - BCP/DRP relationship, 90
 - Special Publication 800-37, certification and accreditation, 15
 - Non-DAC (Nondiscretionary Access Control), 22
 - Nondisclosure Agreement, *see* NDA
 - Nondiscretionary Access Control, *see* Non-DAC
 - Nonrepudiation
 - cryptography, 40
 - definition, 21
 - Normalization
 - databases, 141
 - forms, 141
 - NSA (National Security Agency), Orange Book, 83
 - NVR (Network Video Recorder), and CCTV, 56
- O**
- Object (“black box”) and OOP, 134, 135
 - definition, 134
 - example, 134*f*
 - Object (data file), 21
 - Object Linking and Embedding, *see* OLE
 - Object Management Group, *see* OMG
 - Object Request Broker, *see* ORB
 - Object-Oriented Programming, *see* OOP
 - OCTAVE (Operationally Critical Threat, Asset, and Vulnerability Evaluation), 13
 - OECD (Organization for Economic Cooperation and Development), privacy guidelines, 169
 - OEP (Occupant Emergency Plan) as BCP plan, 98
 - BCP–DRP relationship, 90
 - definition, 99*t*
 - OFB (Output Feedback), 43, 43*t*
 - OFDM (Orthogonal Frequency-Division Multiplexing), 122
 - Offshoring, 13
 - OLE (Object Linking and Embedding), 136
 - OMG (Object Management Group), and CORBA, 136
 - One-time passwords, as Type 1 authentication, 27
 - OOP (Object-Oriented Programming)
 - cornerstone concepts, 134–135
 - ORBs, 136
 - overview, 134–136
 - Open Source Interconnection Reference Model, *see* OSI Reference Model
 - Open system, 71
 - Open-source software, 130
 - Operationally Critical Threat, Asset, and Vulnerability Evaluation, *see* OCTAVE
 - Operations security
 - administrative security, 147–149
 - computer/network attacks, 155–156
 - continuity of operations, 150–155
 - definition, 147
 - sensitive information/media security, 149–150
 - Orange Book, *see* TCSEC
 - ORB (Object Request Broker), 136
 - Organization for Economic Cooperation and Development, *see* OECD
 - Organizationally Unique Identifier, *see* OUI
 - Orthogonal Frequency-Division Multiplexing, *see* OFDM
 - OSI (Open Source Interconnection) Reference Model
 - definition, 110–112
 - example, 111*t*
 - vs.* TCP/IP, 112*t*, 113
 - OUI (Organizationally Unique Identifier), 113
 - Output Feedback, *see* OFB
 - Outsourcing, 13
 - Overt channel, 78
 - Overwriting, *see* Wiping
- P**
- Packet filter, 119
 - Packet-switched networks, 110
 - PACs (Privilege Attribute Certificates), SESAME, 33
 - PAN (Personal Area Network), 110
 - PAP (Password Authentication Protocol), 23–24, 120–121
 - Paper shredders, 61
 - Parallel processing, DRP, 101
 - Parity, 152
 - Partial knowledge test, 34
 - Passive Infrared sensor, *see* PIR sensor
 - Passphrases, as Type 1 authentication, 27
 - Password Authentication Protocol, *see* PAP
 - Password cracking, 156
 - Password guessing, 156
 - Passwords, as Type 1 authentication, 27
 - Patent, legal issues, 165
 - PATRIOT Act of 2001
 - and Computer Fraud and Abuse Act, 167
 - overview, 168
 - PCI (Payment Card Industry), as control framework, 13
 - PCI-DSS (Payment Card Industry Data Security Standard), 85
 - PEAP (Protected EAP), 121
 - Penetration tests, 33–34
 - Perimeter defenses
 - bollards, 56
 - CCTV, 56
 - contraband checks, 58
 - definition, 55–59

- Perimeter defenses (*Continued*)
- dogs, 59
 - doors/windows, 58
 - fences, 55
 - gates, 55
 - guards, 59
 - lights, 56
 - locks, 57
 - mantraps/turnstiles, 58
 - motion detectors and alarms, 58
 - smart cards, 57–58
 - tailgating, 58
 - walls/floors/ceilings, 59
- Permanent Virtual Circuit, *see* PVC
- Permutation, in cryptography, 40
- Personal Area Network, *see* PAN
- Personal Identification Number, *see* PIN
- Personally Identifiable Information, *see* PII
- PGP (Pretty Good Privacy), 52
- Photoelectric sensors, environmental controls, 63
- Physical controls, 25
- Physical destruction, media security, 150
- Physical (environmental) security
- definition, 55
 - environmental controls, 61–64
 - perimeter defenses, 55–59
 - site design/configuration, 60
 - site selection issues, 59–60
 - system defenses, 60–61
- Physical layer, 111
- Piggybacking, *see* Tailgating
- PII (Personally Identifiable Information)
- confidentiality, 19
 - covert channels, 77
 - privacy, 12–13
- PIN (Personal Identification Number), as Type 1 authentication, 26
- Ping of Death, 157*t*
- Pipelining, 72
- PIR (Passive Infrared) sensor, as perimeter defense, 58
- PKI (Public Key Infrastructure), 50
- Plain Old Telephone Service, *see* POTS
- Plaintext
- cryptographic attacks, 47
 - definition, 39
 - strong encryption, 40
- PLD (Programmable Logic Device), 76
- Point-to-Point Protocol, *see* PPP
- Policy
- ACLs, 25
 - basic considerations, 9–11
 - and covert channels, 77
 - definition, 9–10
 - example, 11*t*
 - ITSEC, 84
 - kernel enforcement, 77
 - program components, 9–10
 - types, 10
- Polyalphabetic cipher, 40
- Polyinstantiation
- definition, 80
 - and OOP, 134, 135, 135*f*
- Polymorphic virus, 79
- Polymorphism, and OOP, 134, 135, 135*f*
- POP (Post Office Protocol), 115
- Port controls, system defenses, 61
- Portable fire extinguishers, for fire suppression, 66
- POST (Power-On Self-Test), 76
- POTS (Plain Old Telephone Service), 124
- PP (Protection Profile), 84
- PPP (Point-to-Point Protocol), 122
- Pre-action systems, for fire suppression, 66
- Presentation layer, 112
- Pretty Good Privacy, *see* PGP
- Preventive controls
- definition, 25
 - identification, 97
- Prime number factoring, method, 45
- Privacy, 12–13, 169–170
- Privacy Act of 1974, 169
- Private key, asymmetric encryption, 45
- Privilege Attribute Certificates, *see* PACs
- Procedure
- and Access Control, 24–25
 - definition, 10
 - example, 11*t*
- Process
- definition, 73
 - isolation, 75
- Programmable Logic Device, *see* PLD
- Programming concepts
- bytecode, 130
 - compilers, 130
 - interpreters, 130
 - machine code, 129–130
 - publicly released software, 130
- PROM (Programmable Read Only Memory), 76
- “Protect society, the commonwealth, and the infrastructure,” 171
- Protected EAP, *see* PEAP
- Protection Profile, *see* PP
- “Provide diligent and competent service to principals,” 171
- Proxy firewalls
- application-layer, 119, 120
 - circuit-level, 120
 - definition, 119–120
- Pseudo guard, 59
- Public key, asymmetric encryption, 45
- Public Key Infrastructure, *see* PKI
- Publicly released software, 130
- Punitive damages, 163*t*
- Purpose, security policy definition, 9
- PVC (Permanent Virtual Circuit), 117
- ## Q
- Qualitative Risk Analysis
- definition, 7–8
 - via Risk Analysis Matrix, 3–4
- Quantitative Risk Analysis, 7–8
- Query languages
- databases, 139, 142
 - injection, 137
- ## R
- Race conditions, *see* TOCTOU
- RAD (Rapid Application Development), 133
- Radio-Frequency Identification, *see* RFID
- RADIUS (Remote Authentication Dial in User Service), 23
- RAID (Redundant Array of Inexpensive Disks)
- as fault tolerance, 151–154
 - levels, 154*t*
 - mirroring/stripping/parity, 152
 - reliability, 153
- RAID 0
- definition, 152
 - example, 152*f*
 - levels, 154*t*
- RAID 1
- definition, 152
 - example, 152*f*
 - levels, 154*t*
- RAID 1+0, 153–154
- RAID 2, 153
- RAID 3
- definition, 153
 - levels, 154*t*

- RAID 4
 - definition, 153
 - levels, 154t
 - RAID 5
 - definition, 153
 - example, 154f
 - levels, 154t
 - RAID 6
 - definition, 153
 - levels, 154t
 - RAM (Random Access Memory), 73–74
 - Rapid Application Development, *see* RAD
 - RBAC (Role-Based Access Controls)
 - Access control procedures, 24
 - definition, 22
 - RC4 (Rivest Cipher 4), 124
 - RC5 (Rivest Cipher 5), 44–45
 - RC6 (Rivest Cipher 6), 44–45
 - Read Only Memory, *see* ROM
 - Real evidence, 167
 - Reciprocal agreement, 98
 - Reconstitution, disaster recovery
 - process, 93
 - Recovery controls, 26
 - Recovery metrics, BIA, 95–96
 - Recovery Point Objective, *see* RPO
 - Recovery strategy
 - cold site, 98
 - hot site, 97
 - mobile site, 98
 - and MTD, 97–98
 - reciprocal agreement, 98
 - redundant site, 97
 - warm site, 97–98
 - Recovery Time Objective, *see* RTO
 - Red Book, *see* TNI
 - Reduced Instruction Set Computer, *see* RISC
 - Redundant Array of Inexpensive Disks, *see* RAID
 - Redundant site, 97
 - Redundant systems, 155
 - Reference monitor, 77
 - Referential integrity
 - definition, 140–141
 - and foreign keys, 141
 - Relational databases
 - definition, 139–142
 - example table, 140t
 - Religious law, as legal system, 162
 - Remote access, secure
 - communications, 124–125
 - Remote Authentication Dial in User Service, *see* RADIUS
 - Remote File Inclusion, *see* RFI
 - Remote Procedure Calls, *see* RPCs
 - Repeaters, 117
 - Response step, disaster recovery
 - process, 92
 - Retina scan
 - biometric controls, 29
 - risks, 30
 - Return On Investment, *see* ROI
 - Review, DRP testing, 100
 - RFI (Remote File Inclusion), 137
 - RFID (Radio-Frequency Identification), smart cards, 57
 - Ring model
 - definition, 70
 - example, 70f
 - usage, 70b
 - RISC (Reduced Instruction Set Computer), 73
 - Risk acceptance
 - criteria, 7
 - as risk choice, 7
 - Risk analysis
 - ALE calculation, 4–5
 - assets, 2
 - basic concept, 1–9
 - equations summary, 5t
 - impact, 3
 - and outsourcing/offshoring, 13
 - qualitative and quantitative, 7–8
 - risk choices, 6–7
 - risk formula, 3
 - risk management process, 8–9
 - TCO, 5
 - threats and vulnerabilities, 2
 - Risk Analysis Matrix
 - example, 3t
 - function, 3–4
 - Qualitative Risk Analysis as, 8
 - Risk assessment, BCP/DRP-focused, 95
 - Risk avoidance, as risk choice, 7
 - Risk management
 - overview, 1
 - process, 8–9
 - Risk mitigation
 - BIA, 97
 - definition, 7
 - Risk transfer, 7
 - Rivest Cipher 4, *see* RC4
 - Rivest Cipher 5, *see* RC5
 - Rivest Cipher 6, *see* RC6
 - Robust Security Network, *see* RSN
 - ROI (Return On Investment), 5–6
 - Role-Based Access Controls, *see* RBAC
 - ROM (Read Only Memory), 74
 - Rootkit, 79, 157t
 - Rotation of duties, 148
 - Routers, 118
 - RPCs (Remote Procedure Calls), 111
 - RPO (Recovery Point Objective), 96
 - RSN (Robust Security Network), 124
 - RTO (Recovery Time Objective), 96
- ## S
- S/MIME (Secure/Multipurpose Internet Mail Extensions), 52
 - SA (Security Association), 51
 - Safeguards, 21
 - SAML (Security Assertion Markup Language), 79
 - Sanction
 - and accountability, 21
 - and compliance, 10
 - SANS Institute, 137
 - Sarbanes-Oxley, *see* SOX (Sarbanes-Oxley)
 - Scope, security policy definition, 9
 - SCP (Secure Copy), and SSH, 115
 - SDES (Single Data Encryption Standard), 43
 - SDLC (Applications), 133
 - SDSL (Single-Line Digital Subscriber Line)
 - definition, 125
 - speed and distance, 125t
 - Second session key, *see* C/S session key
 - “Secret key” encryption, 41
 - Secure communications
 - authentication protocols and frameworks, 120–121
 - basic concept, 120–125
 - remote access, 124–125
 - VPN, 121–122
 - WLANs, 122–124
 - Secure Copy, *see* SCP
 - Secure European System for Applications in a Multi-vendor Environment, *see* SESAME
 - Secure hardware architecture
 - basic considerations, 71–76
 - computer bus, 71, 71f
 - CPU, 71–73
 - memory, 73–74
 - memory protection, 74–76
 - system unit/motherboard, 71

- Secure Hash Algorithm, *see* SHA
- Secure Hash Algorithm 1, *see* SHA-1
- Secure Multipurpose Mail Extensions, *see* S/MIME
- Secure operating systems, architecture, 76–77
- Secure Shell, *see* SSH
- Secure Sockets Layer, *see* SSL
- Secure system design
 - abstraction, 70
 - basic considerations, 69–71
 - layering, 69, 70*b*
 - open/closed systems, 71
 - ring model, 70, 70*b*, 70*f*
 - security domain, 70
- Secure/Multipurpose Internet Mail Extensions, *see* S/MIME
- Security Architecture and Design
 - basic considerations, 69
 - operating systems/software, 76–77
 - secure hardware architecture, 71–76
 - secure system design, 69–71
 - security models, 80–83
 - System Vulnerabilities, Threats, and Countermeasures, 77–80
- Security Assertion Markup Language, *see* SAML
- Security assessments, 34
- Security Association, *see* SA
- Security audit, 34
- Security awareness, 12
- Security domain, 70
- Security models
 - access control matrix, 82–83, 83*t*
 - basic considerations, 80–83
 - Bell-LaPadula model, 80
 - Chinese Wall model, 82
 - integrity models, 81–82
 - lattice-based access controls, 80–81, 81*f*
- Security policy, *see* Policy
- Security Property, Bell-LaPadula model, 81
- Security responsibilities
 - examples, 12
 - policy definition, 10
- Security roles, examples, 12
- Security Target, *see* ST (Security Target)
- Security training, 12
- Semantic integrity
 - and attributes, 141
 - definition, 140–141
- Senior Management, role and responsibility, 12
- Sensitive information, security, 149–150
- Separation of duties, 147
- Serial Line Internet Protocol, *see* SLIP
- Service Level Agreement, *see* SLA
- Servicemark, 165
- SESAME (Secure European System for Applications in a Multi-vendor Environment), 33
- Session hijacking, 156
- Session layer, 111
- SFTP, and SSH, 115
- SHA (Secure Hash Algorithm), 46
- SHA-1 (Secure Hash Algorithm 1)
 - digital signatures, 48, 49, 49*f*
 - and evidence integrity, 167
 - hash functions, 46
 - HMACs, 50
 - and IKE, 51
 - and VPN, 121–122
- Shareware, 130
- Shielded Twisted Pair, *see* STP
- Shredding, media security, 150
- Side-channel attacks, 48
- Simple Integrity Axiom, Biba model, 82
- Simple Mail Transfer Protocol, *see* SMTP
- Simple Security Property, Bell-LaPadula model, 81
- Simplex communication, 109
- Simulation test, DRP, 101
- Single Data Encryption Standard, *see* SDES
- Single Loss Expectancy, *see* SLE
- Single Sign-On, *see* SSO
- Single-Line DSL, *see* SDSL
- Site configuration, as physical security, 60
- Site design, as physical security, 60
- Site selection, associated issues, 59–60
- SLA (Service Level Agreement), 150
- Slack space, 166
- SLE (Single Loss Expectancy)
 - definition, 4
 - encrypted laptop ALE, 6*t*
 - risk equations, 5*t*
 - unencrypted laptop ALE, 6*t*
- SLIP (Serial Line Internet Protocol), 122
- Smart card, 57–58
- SMDS (Switched Multimegabit Data Service), 117
- Smoke detectors, environmental controls, 63
- SMP (Symmetric Multiprocessing), 73
- SMTP (Simple Mail Transfer Protocol), 115
- Smurf attack, 157*t*
- Sniffer, 118
- Social engineering, 34
- SOCKS, 120
- Soda acid, as fire suppression agent, 64
- Software architecture, basic considerations, 76–77
- Software licenses, legal issues, 165
- Software vulnerabilities
 - disclosure, 138
 - types, 137–138
- Someplace you are, biometric controls, 31
- Something you are, *see* Type 3 authentication
- Something you have, *see* Type 2 authentication
- Something you know, *see* Type 1 authentication
- Source code, 129–130
- SOX (Sarbanes-Oxley), and offshoring, 13
- Spiral Model
 - definition, 131
 - example, 132*f*
- Spring-bolt lock, 57
- Sprinkler systems
 - example, 66*f*
 - for fire suppression, 65–66
- SQL (Structured Query Language)
 - databases, 139
 - injection, 137
- SRAM (Static Random Access Memory), 74
- SSH (Secure Shell)
 - application layer, 113
 - definition, 115
- SSL (Secure Sockets Layer)
 - definition, 51–52
 - secure communications, 122
- SSO (Single Sign-On)
 - via centralized access control, 22
 - definition, 31
- ST (Security Target), 84
- Standard
 - definition, 11
 - example, 11*t*
- Stateful firewalls, 119
- Static environmental controls, 63

- Static passwords, as Type 1 authentication, 27
 - Static Random Access Memory, *see* SRAM
 - Statutory damages, 163*t*
 - Stealth virus, 79
 - Storage
 - media, 61
 - media security, 149
 - Storage channel, 78
 - STP (Shielded Twisted Pair), and EMI, 62
 - Stream ciphers, 41
 - Striping, 152
 - Strong authentication, 27
 - Strong Tranquility Property, Bell-LaPadula model, 81
 - Structured Query Language, *see* SQL
 - Structured walkthrough, DRP testing, 100
 - Subject, 21
 - Substitution, in cryptography, 40
 - Surge protectors, environmental controls, 62
 - SVC (Switched Virtual Circuit), 117
 - Swapping
 - definition, 75
 - virtual memory, 75
 - Switched Multimegabit Data Service, *see* SMDS
 - Switches
 - definition, 118
 - example, 119*f*
 - Symmetric Encryption
 - AES, 44, 44*t*
 - vs.* asymmetric, 46
 - basic concept, 41–45
 - Blowfish/Twofish, 44
 - definition, 41
 - DES, 41–44
 - International Data Encryption Algorithm, 44
 - stream and block ciphers, 41
 - Symmetric Processing, *see* SMP
 - SYN Flood, 157*t*
 - Synchronous Dynamic Tokens, as Type 2 authentication, 27
 - System defenses and types, 60–61
 - System integrity, 20
 - System redundancy, 154–155
 - System unit, 71
 - System Vulnerabilities, Threats, and Countermeasures
 - backdoors, 78
 - covert channels, 77–78
 - database security, 80
 - definition, 77–80
 - malware, 78–79
 - TOCTOU, 78
 - Systems Development Lifecycle, *see* SDLC (applications)
- T**
- T1, 116
 - T3, 116
 - Tabletop exercise, DRP testing, 100
 - TACACS (Terminal Access Controller Access Control System), 23
 - TACACS+ (Terminal Access Controller Access Control System), 23
 - Tailgating, 58
 - Tape encryption, system defenses, 61
 - Target of Evaluation, *see* ToE
 - TCO (Total Cost of Ownership)
 - definition, 5
 - and outsourcing, 13
 - TCP (Transmission Control Protocol)
 - definition, 111
 - and DNS, 115
 - function, 114
 - ports, 114
 - TCP/IP (Transmission Control Protocol/Internet Protocol)
 - application-layer protocols, 115–116
 - definition, 112–113
 - vs.* OSI, 112*t*, 113
 - TCSEC (Trusted Computer System Evaluation Criteria)
 - characteristics, 83
 - divisions, 83
 - TDEA (Triple Data Encryption Algorithm), *see* TDES
 - TDES (Triple Data Encryption Standard)
 - cryptographic strength, 40
 - definition, 43–44
 - encryption order/keying options, 44
 - Team activation, disaster recovery process, 92–93
 - Teardrop, 157*t*
 - Technical controls, 25
 - Telecommunications
 - network architecture and design, 109–117
 - network devices, 117–120
 - secure communications, 120–125
 - Telnet
 - application layer, 113
 - definition, 115
 - Temporal Key Integrity Protocol, *see* TKIP
 - Ten Commandments of Computer Ethics, 170
 - Tenancy, site configuration issues, 60
 - Terminal Access Controller Access Control System, *see* TACACS (Terminal Access Controller Access Control System); TACACS+
 - TFTP (Trivial File Transfer Protocol), 115
 - TGT (Ticket Granting Ticket), Kerberos steps, 32
 - Thin client, 77
 - Thread, 73
 - Threat
 - definition, 2
 - disasters, 90, 91
 - disruptive events, 90, 92, 92*t*
 - risk formula, 3
 - Throughput, 28
 - Ticket Granting Ticket, *see* TGT
 - Time of Check/Time of Use, *see* TOCTOU
 - TKIP (Temporal Key Integrity Protocol), WPA2, 124
 - TLS (Transport Layer Security)
 - definition, 51–52
 - secure communications, 122
 - TNI (Trusted Network Interpretation), 83
 - TOCTOU (Time of Check/Time of Use), 78
 - ToE (Target of Evaluation), 84
 - Topography, site selection issues, 60
 - Total Cost of Ownership, *see* TCO
 - Trade secret, legal issues, 165
 - Trademark, legal issues, 165
 - Training
 - Call Tree, 101
 - DRP, 101
 - Transmission Control Protocol, *see* TCP
 - Transmission Control Protocol/Internet Protocol, *see* TCP/IP
 - Transport layer
 - definition, 111
 - protocols, 113–115

Transport Layer Security, *see* TLS
 Transport mode, 51
 Transposition, *see* Permutation
 Triple Data Encryption Standard, *see* TDES
 Trivial File Transfer Protocol, *see* TFTP ()
 Trojan horse, 78–79, 157*t*
 Trusted Computer System Evaluation Criteria, *see* TCSEC
 Trusted Network Interpretation, *see* TNI
 Truth table, XOR, 41*t*
 Tunnel mode, 51
 Tuple
 definition, 139–140
 and entity integrity, 141
 Turnstile, 58
 Twofish, 44
 Type 1 authentication, 26–27
 Type 2 authentication, 27–28
 Type 3 authentication, 28–30
 Type I error, 28
 Type II error, 28

U

UDP (User Datagram Protocol)
 definition, 111
 and DNS, 115
 function, 114
 Unallocated space, 166
 UPSs (Uninterruptible Power Supplies), environmental controls, 62
 User, role and responsibility, 12
 Utility reliability, site selection issues, 60
 UTP (Unshielded Twisted Pair)
 baseband, 110
 and EMI, 62

V

VDSL (Very High Rate Digital Subscriber Line)
 definition, 125
 speed and distance, 125*t*
 VHS, and CCTV, 56
 Views, databases, 142, 142*t*
 Virtual memory, 75
 Virtualization, 77
 Virus
 definition, 78–79, 157*t*
 types, 78*b*
 Voice print, biometric controls, 30
 VPN (Virtual Private Network)
 definition, 50
 secure communications, 121–122
 Vulnerability
 definition, 2
 risk formula, 3
 software, 137–138
 Vulnerability scanning, 34

W

Walkthrough drill, DRP, 101
 Walls
 fire rating, 59
 as perimeter defense, 59
 WAN (Wide Area Network)
 definition, 110
 technologies and protocols, 116–117
 War dialing, 34
 Warded lock, 57
 Warm site, 97–98
 Water, as fire suppression agent, 64
 Waterfall Model, 131
 Weak Tranquility Property, Bell-LaPadula model, 81
 WEP (Wired Equivalent Privacy), 123–124

Wet chemical, as fire suppression agent, 65
 Whole-disk encryption, system defenses, 61
 Wide Area Network, *see* WAN
 Wi-Fi Protected Access, *see* WPA
 Wi-Fi Protected Access 2, *see* WPA2
 Windows, as perimeter defense, 58
 Wiping, media security, 150
 Wire Local Area Networks, *see* WLANs
 Wired Equivalent Privacy, *see* WEP
 WLANs (Wire Local Area Networks)
 802.11, 123
 802.11i, 124
 definition, 122–124
 FHSS/DSSS/OFDM, 123
 WEP, 123–124
 WORM (Write Once Read Many), 76
 Worm, 78–79, 157*t*
 WPA (Wi-Fi Protected Access), 124
 WPA2 (Wi-Fi Protected Access 2), 124
 WRT (Work Recovery Time), 96

X

X.25, 117
 XML (eXtensible Markup Language), 79
 XOR (Exclusive Or)
 and CBC, 42
 definition, 40–41
 and OFB, 43
 truth table, 41*t*
 XP (eXtreme Programming), 133
 XSRF, *see* CSRF
 XSS (Cross-Site Scripting), 137

Z

Zero knowledge test, 34
 Zero-day exploits, 78