

Second Edition



Certified in Cybersecurity

STUDY GUIDE

**COVERS 2026-2029 CC EXAM OBJECTIVES
AND ISC2 EXAM GUIDANCE FOR AI**

Includes interactive online learning environment and study tools with:

A full complete practice exam

Electronic flash cards

Searchable key term glossary

MIKE CHAPPLE, Ph.D.
CC, CISSP, CCSP

 **SYBEX**
A Wiley Brand

Table of Contents

[Cover](#)

[Table of Contents](#)

[Title Page](#)

[Copyright](#)

[Acknowledgments](#)

[About the Author](#)

[About the Technical Editor](#)

[Introduction](#)

[Part I: Domain 1: Security Principles](#)

[Chapter 1: Confidentiality, Integrity, Availability, and Non-repudiation](#)

[The CIA Triad](#)

[Non-repudiation](#)

[Chapter 2: Authentication, Authorization, and Accounting \(AAA\)](#)

[Access Control Process](#)

[Password Policies](#)

[Authentication Factors](#)

[Chapter 3: Privacy](#)

[Privacy](#)

[Privacy Management Framework](#)

[Chapter 4: Risk Management](#)

[Risk Categories and Types](#)

[Risk Identification and Assessment](#)

[Risk Treatment Strategies](#)

[Risk Profile and Tolerance](#)

[Chapter 5: Governance Concepts](#)

[Laws and Regulations](#)

[Frameworks and Guidelines](#)

[Security Policies, Standards, and Procedures](#)

[Chapter 6: Cybersecurity Controls](#)

[What Are Cybersecurity Controls?](#)

[Categorizing Security Controls](#)

[Chapter 7: Professional and Ethical Conduct](#)

[Corporate Codes of Ethics](#)

[Due Care and Due Diligence](#)

[ISC2 Code of Ethics](#)

[Ethical Complaint Procedures](#)

[Code of Professional Conduct](#)

[Part II: Domain 2: Security Governance](#)

[Chapter 8: Governance, Risk, and Compliance](#)

[Purpose of Governance, Risk, and Compliance](#)

[Importance of GRC](#)

[Frameworks and Tools](#)

[Chapter 9: Business Continuity](#)

[Business Continuity Planning](#)

[Business Continuity Controls](#)

[High Availability and Fault Tolerance](#)

[Chapter 10: Disaster Recovery](#)

[Disaster Recovery Planning](#)

[Backups](#)

[Disaster Recovery Sites](#)

[Testing Disaster Recovery Plans](#)

[Chapter 11: Security Awareness](#)

[Organizational Culture](#)

[Social Engineering](#)

[Phishing](#)

[Password Protection](#)

[Security Education](#)

[Chapter 12: Measuring Cybersecurity Effectiveness](#)

[Cybersecurity Metrics](#)

[Dashboards, Scorecards, and Reports](#)

[Part III: Domain 3: Identity and Access Management \(IAM\) Concepts](#)

[Chapter 13: Identity Life Cycle Management](#)

[The Identity Life Cycle](#)

[Identity Management Frameworks and Tools](#)

[Chapter 14: Logical Access Controls](#)

[Authorization](#)

[Part IV: Domain 4: Networking and Cloud Security Concepts](#)

[Chapter 15: Network Security](#)

[Network Types](#)

[TCP/IP Networking](#)

[IP Addressing](#)

[Securing Wi-Fi Networks](#)

[VPN](#)

[Embedded Systems](#)

[IoT](#)

[Chapter 16: Network Security Architecture](#)

[Defense in Depth](#)

[Zero Trust](#)

[Network Segmentation](#)

[Chapter 17: Cloud Security](#)

[Cloud Characteristics](#)

[Cloud Deployment Models](#)

[Cloud Service Models](#)

[Security and the Shared Responsibility Model](#)

[Part V: Domain 5: Security Operations and Incident Response](#)

[Chapter 18: Encryption](#)

[Cryptography](#)

[Encryption Algorithms](#)

[Uses of Encryption](#)

[Hash Functions](#)

[Quantum-Resistant Cryptography](#)

[Chapter 19: Data Handling](#)

[Data Life Cycle](#)

[Data Masking](#)

[Data Classification](#)

[Chapter 20: Security Operations](#)

[Logging](#)

[Log Monitoring](#)

[Security Event Triage](#)

[Threat Actors](#)

[Cyber Threat Intelligence](#)

[Threat Frameworks](#)

[Chapter 21: Incident Response](#)

[Creating an Incident Response Program](#)

[Building an Incident Response Team](#)

[Incident Response Exercises](#)

[Incident Communications Plan](#)

[Incident Identification and Response](#)

[Chapter 22: Asset Protection](#)

[Asset Lifecycle Management](#)

[Configuration Management](#)

[Change Management](#)

[Chapter 23: Security Testing](#)

[Security Readiness Testing](#)

[Application Testing](#)

[Physical Penetration Testing](#)

[Chapter 24: Artificial Intelligence](#)

[What Is Artificial Intelligence?](#)

[AI Security Principles](#)

[AI and Organizational Resilience](#)

[AI and Access Controls](#)

[AI and Network Security](#)

[AI and Security Operations](#)

[Index](#)

[Advertisement Page: Get Certified!](#)

[Advertisement Page: Online Test Bank](#)

[End User License Agreement](#)

List of Illustrations

Chapter 1

[Figure 1.1 The CIA triad summarizes the three key goals of information security: confidentiality...](#)

Chapter 2

[Figure 2.1 The physical access control process.](#)

[Figure 2.2 The digital access control process.](#)

[Figure 2.3 Creating a password in LastPass.](#)

[Figure 2.4 Fingerprint authentication on a smartphone.](#)

[Figure 2.5 Iris scan authentication for entering a facility.](#)

Chapter 3

[Figure 3.1 Do you have a reasonable expectation of privacy?](#)

Chapter 4

[Figure 4.1 The risk management lifecycle.](#)

[Figure 4.2 Risks are the combination of a threat and a corresponding vulnerability.](#)

[Figure 4.3 Qualitative risk assessment.](#)

[Figure 4.4 Applying controls reduces the inherent risk down to the residual risk but also i...](#)

Chapter 5

[Figure 5.1 NIST Cybersecurity Framework \(CSF\) core functions.](#)

[Figure 5.2 NIST risk management framework \(RMF\).](#)

Chapter 7

[Figure 7.1 AT&T's code of business conduct.](#)

Chapter 9

[Figure 9.1 Web-based application.](#)

[Figure 9.2 Adding clustered web servers.](#)

[Figure 9.3 Adding high availability firewalls.](#)

[Figure 9.4 Adding redundant network links.](#)

[Figure 9.5 Server with dual power supplies.](#)

[Figure 9.6 Uninterruptible power supply \(UPS\).](#)

[Figure 9.7 RAID 1 disk mirroring.](#)

[Figure 9.8 RAID 5 disk striping with parity.](#)

Chapter 10

[Figure 10.1 LTO backup tapes.](#)

Chapter 11

[Figure 11.1 A security awareness poster.](#)

Chapter 13

[Figure 13.1 The identity life cycle.](#)

Chapter 14

[Figure 14.1 A Microsoft Windows access control list.](#)

Chapter 15

[Figure 15.1 The open systems interconnection \(OSI\) model.](#)

[Figure 15.2 Communication from a user to a web server.](#)

[Figure 15.3 Communication from a web server to a user.](#)

[Figure 15.4 SSIDs appearing on a macOS system.](#)

[Figure 15.5 A captive portal used for Wi-Fi authentication.](#)

[Figure 15.6 Wi-Fi standards summary.](#)

Chapter 16

[Figure 16.1 A network border firewall with three interfaces.](#)

[Figure 16.2 A typical network diagram.](#)

[Figure 16.3 A typical VLAN layout.](#)

Chapter 18

[Figure 18.1 A plaintext message.](#)

[Figure 18.2 The ciphertext message obtained by encrypting the plaintext in Figure 18.1.](#)

[Figure 18.3 Examples of symmetric shapes.](#)

[Figure 18.4 Symmetric encryption with two individuals.](#)

[Figure 18.5 Symmetric encryption with three individuals.](#)

[Figure 18.6 Symmetric encryption with more participants.](#)

Chapter 19

[Figure 19.1 Data life cycle.](#)

[Figure 19.2 NIST storage sanitization flowchart.](#)

Chapter 20

[Figure 20.1 The ATT&CK definition for active scanning.](#)

[Figure 20.2 The Cyber Kill Chain.](#)

[Figure 20.3 A Diamond Model analysis of a compromised system.](#)

Chapter 21

[Figure 21.1 NIST incident response life cycle.](#)

Chapter 22

[Figure 22.1 The asset management lifecycle.](#)

[Figure 22.2 The change management process.](#)

List of Tables

Chapter 9

[Table 9.1 Example of Prioritized Risks and Potential Impacts](#)

CCSM Certified in Cybersecurity Study Guide

Second Edition

Mike Chapple, Ph.D.
CC, CISSP, CCSP



Copyright © 2026 by John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies.

Published by John Wiley & Sons, Inc., Hoboken, New Jersey.

No part of this publication may be reproduced, stored in a retrieval system, or transmitted in any form or by any means, electronic, mechanical, photocopying, recording, scanning, or otherwise, except as permitted under Section 107 or 108 of the 1976 United States Copyright Act, without either the prior written permission of the Publisher, or authorization through payment of the appropriate per-copy fee to the Copyright Clearance Center, Inc., 222 Rosewood Drive, Danvers, MA 01923, (978) 750-8400, fax (978) 750-4470, or on the web at www.copyright.com. Requests to the Publisher for permission should be addressed to the Permissions Department, John Wiley & Sons, Inc., 111 River Street, Hoboken, NJ 07030, (201) 748-6011, fax (201) 748-6008, or online at <http://www.wiley.com/go/permission>.

The manufacturer's authorized representative according to the EU General Product Safety Regulation is Wiley-VCH GmbH, Boschstr. 12, 69469 Weinheim, Germany, e-mail: Product_Safety@wiley.com.

Trademarks: Wiley and the Wiley logo, and the Sybex logo, are trademarks or registered trademarks of John Wiley & Sons, Inc. and/or its affiliates in the United States and other countries and may not be used without written permission. All other trademarks are the property of their respective owners. John Wiley & Sons, Inc. is not associated with any product or vendor mentioned in this book.

Limit of Liability/Disclaimer of Warranty: While the publisher and the authors have used their best efforts in preparing this work, including a review of the content of the work, neither the publisher nor the authors make any representations or warranties with respect to the accuracy or completeness of the contents of this work and specifically disclaim all warranties, including without limitation any implied warranties of merchantability or fitness for a particular purpose. Certain AI systems have been used in the creation of this work. No warranty may be created or extended by sales representatives, written sales materials or promotional statements for this work. The fact that an organization, website, or product is referred to in this work as a citation and/or potential source of further information does not mean that the publisher and authors endorse the information or services the organization, website, or product may provide or recommendations it may make. This work is sold with the understanding that the publisher is not engaged in rendering professional services. The advice and strategies contained herein may not be suitable for your situation. You should consult with a specialist where appropriate. Further, readers should be aware that websites listed in this work may have changed or disappeared between when this work was written and when it is read. Neither the publisher nor authors shall be liable for any loss of profit or any other commercial damages, including but not limited to special, incidental, consequential, or other damages.

For general information on our other products and services or for technical support, please contact our Customer Care Department within the United States at (800) 762-2974, outside the United States at (317) 572-3993 or fax (317) 572-4002. For product technical support, you can find answers to frequently asked questions or reach us via live chat at <https://sybexsupport.wiley.com>.

Wiley also publishes its books in a variety of electronic formats. Some content that appears in print may not be available in electronic formats. For more information about Wiley products, visit our web site at www.wiley.com.

Library of Congress Cataloging-in-Publication Data has been applied for:

ISBN: 9781394454891 (Paperback)

ISBN: 9781394454907 (epub)

ISBN: 9781394454914 (epdf)

Cover Design: Wiley

Cover Image: Jeremy Woodhouse/Getty Images

Acknowledgments

Books like this involve work from many people, and as an author, I truly appreciate the hard work and dedication that the team at Wiley shows. I would especially like to thank my acquisitions editor, Jim Minatel. I've worked with Jim for too many years to count, and it's always an absolute pleasure working with a true industry pro.

I also greatly appreciate the editing and production team for the book, including Tracy Brown, the project editor, who brought years of experience and great talent to the project; and Shahla Pirnia, the technical editor, who provided insightful advice and gave wonderful feedback throughout the book. I would also like to thank the many behind-the-scenes contributors, including the graphics, production, and technical teams who make the book and companion materials into a finished product.

My agent, Carole Jelen of Waterside Productions, continues to provide me with wonderful opportunities, advice, and assistance throughout my writing career.

Finally, I would like to thank my family, who supported me through the late evenings, busy weekends, and long hours that a book like this requires to write, edit, and get to press.

About the Author

Mike Chapple, CISSP, CCSP, is an author of the best-selling *CISSP ISC2 Certified Information Systems Security Professional Official Study Guide* (Sybex, 2024), now in its tenth edition. He is an information security professional with 30 years of experience in higher education, the private sector, and government.

Mike currently serves as Teaching Professor of IT, Analytics, and Operations at the University of Notre Dame's Mendoza College of Business. He previously served as Senior Director of IT Service Delivery at Notre Dame, overseeing information security, data governance, IT architecture, project management, strategic planning, and product management for the university.

Before returning to Notre Dame, Mike served as Executive Vice President and Chief Information Officer of the Brand Institute, a Miami-based marketing consultancy. Mike also spent four years in the information security research group at the National Security Agency and served as an active duty intelligence officer in the U.S. Air Force.

Mike is the author of more than 50 books, including *Cyberwarfare: Information Operations in a Connected World* (Jones & Bartlett, 2022), *ISC2 CISSP Official Study Guide* (Wiley, 2024), and *CompTIA Cybersecurity Analyst+ (CySA+) Study Guide* (Wiley, 2026) and *Practice Tests* (Wiley, 2026).

Mike earned both his BS and PhD degrees from Notre Dame in computer science and engineering. He also holds an MS in computer science from the University of Idaho and an MBA from Auburn University. His IT certifications include the CC, CISSP, Security+, CySA+, CISA, PenTest+, CIPP/US, CISM, and CCSP credentials.

Mike provides books, video-based training, and free study groups for a wide variety of IT certifications at his website, CertMike.com.

About the Technical Editor

Shahla Pirnia is a freelance technical editor and proofreader with a focus on cybersecurity and certification topics. She currently serves as a technical editor for [CertMike.com](https://certmike.com). Shahla earned BS degrees in computer and information science and psychology from UMGC and an AA in information systems from Montgomery College, MD. Shahla's certifications include ISC2 CC, CompTIA Security+, Network+, and A+. She holds the Associate of ISC2 designation.

Introduction

If you're preparing to take the Certified in Cybersecurity (CC) exam, you'll undoubtedly want to find as much information as you can about information security. The more information you have at your disposal, the better off you'll be when attempting the exam. This study guide was written with that in mind. The goal is to provide enough information to prepare you for the test, but not so much that you'll be overloaded with information that's outside the scope of the exam.

This book presents the material at an entry level. You don't need any prior experience with cybersecurity to read this book or take the exam. The CC certification is designed for newcomers to the field, and this book will give you all the information you need to know to pass it.

I've included review questions at the end of each chapter to give you a taste of what it's like to take the exam. I recommend that you check out these questions first to gauge your level of expertise. You can then use the book mainly to fill in the gaps in your current knowledge. This study guide will help you round out your knowledge base before tackling the exam.

If you can answer the review questions correctly for a given chapter, you can feel safe moving on to the next chapter. If you're unable to answer them correctly, reread the chapter and try the questions again. Your score should improve.

Note

Don't just study the questions and answers! The questions on the actual exam will be different from the practice questions included in this book. The exam is designed to test your knowledge of a concept or objective, so use this book to learn the objectives behind the questions.

CC Certification

The CC certification is offered by the International Information System Security Certification Consortium, or ISC2, a global nonprofit organization. The mission of ISC2 is to strengthen the influence, diversity, and vitality of the field through advocacy, expertise, and workforce empowerment that accelerates cyber safety and security in an interconnected world. ISC2 achieves this mission by delivering the world's leading information security certification program. The CISSP is the flagship certification in this series and is accompanied by several other ISC2 credentials:

Certified in Cybersecurity (CC)

Systems Security Certified Practitioner (SSCP)

Certified Information Systems Security Professional (CISSP)

Certified Cloud Security Professional (CCSP)

Certified in Governance, Risk, and Compliance (CGRC)

Certified Secure Software Lifecycle Professional (CSSLP)

Information Systems Security Architecture Professional (ISSAP)

Information Systems Security Engineering Professional (ISSEP)

Information Systems Security Management Professional (ISSMP)

The CC certification covers five domains of information security knowledge. These domains are meant to serve as the broad knowledge foundation required to succeed in the information security profession:

Security Principles (24% of exam questions)

Security Governance (17.3% of exam questions)

Identity and Access Management (IAM) Concepts (20% of exam questions)

Networking and Cloud Security Concepts (21.3% of exam questions)

Security Operations and Incident Response (17.3% of exam questions)

Complete details about the CC exam objectives are contained in the Exam Outline. It includes a full outline of exam topics and can be found on the ISC2 website at www.isc2.org/certifications/cc/cc-certification-exam-outline.

Taking the CC Exam

The CC exam uses Computerized Adaptive Testing (CAT) and includes multiple-choice questions and advanced item types. When taking the test, you'll likely find some questions where you think multiple answers might be correct. In those cases, remember that you're looking for the *best* possible answer to the question!

You can find more details about the CC exam and how to take it at www.isc2.org/certifications/cc.

You'll have two hours to take the exam and will be asked to answer 100 to 125 questions. Your exam will be scored on a scale of 1,000 possible points, with a passing score of 700.

Note

The CC exam includes unscored questions that do not count toward your score. ISC2 does this to gather research data, which it then uses when developing new versions of the exam. So, if you come across a question that does not appear to map to any of the exam objectives—or, for that matter, does not appear to belong in the exam—it is likely a seeded question. You never really know whether or not a question is seeded, however, so always make your best effort to answer every question.

Computer-based Testing Environment

The CC exam is administered in a computer-based testing (CBT) format. You can register for the exam through the ISC2 website.

You take the exam in a Pearson VUE testing center located near your home or office. The centers administer many different exams, so you may find yourself sitting in the same room as a student taking a school entrance examination and a health care professional earning a medical certification. If you'd like to become more familiar with the testing environment, the Pearson VUE website offers a virtual tour of a testing center at

<https://www.pearsonvue.com/us/en/test-takers/pearson-professional-center-tour.html>.

When you take the exam, you'll be seated at a computer that has the exam software already loaded and running. It's a pretty straightforward interface that allows you to navigate through the exam. You can launch a demo test from the Pearson VUE website at

<https://www.pearsonvue.com/us/en/test-takers/demo-test.html>.

Exam Tip

At the beginning of the exam, you'll be asked to agree to the terms. This section of the exam has its own three-minute timer. If you don't agree within three minutes, your exam will automatically end and you will not be able to restart it!

Exam Retake Policy

If you don't pass the CC exam, you shouldn't panic. Many individuals don't reach the bar on their first attempt but gain valuable experience that helps them succeed the second time around. When retaking the exam, you'll have the benefit of familiarity with the CBT environment and CC exam format. You'll also have time to study the areas where you felt less confident.

After your first exam attempt, you must wait 30 days before retaking it. If you're not successful on that attempt, you must then wait 60 days before your third attempt and 90 days before your fourth attempt. You can take the exam up to four times within a 12-month period.

Recertification Requirements

Once you've earned your CC credential, you'll need to maintain your certification by paying maintenance fees and participating in continuing professional education (CPE). As long as you maintain your certification in good standing, you will not need to retake the CC exam.

Currently, the annual maintenance fee for the CC credential is \$50 for those who do not hold another ISC2 certification. Members who hold another credential pay a \$125 maintenance fee each year. This fee covers the renewal for all ISC2 certifications held by an individual.

The CC CPE requirement mandates earning at least 45 CPE credits during each three-year renewal cycle. ISC2 provides an online portal where certificate holders can submit CPE completion for review and approval. The portal also tracks annual maintenance fee payments and progress toward recertification.

Using the Online Practice Test

All the chapter practice questions in this book are also available in Sybex's online learning environment, along with an additional online-only full-length 100-question CC practice test. To get access to this online format, go to www.wiley.com/go/sybextestprep and start by registering your book. You'll receive a PIN code and instructions on where to create an online account. Once you have access, you can use the online version to practice in a timed and graded setting.

In addition to the chapter practice questions and practice test, the Sybex online learning environment includes an extensive set of electronic flashcards to improve your exam preparation. Each flashcard has one question and one correct answer. These are great as last-minute drills. And there is an online glossary with a searchable list of key terms introduced in this study guide that you should know for the CC certification exam.

How to Contact the Publisher

If you believe you have found a mistake in this book, please bring it to our attention. At John Wiley & Sons, we understand how important it is to provide our customers with accurate content, but even with our best efforts an error may occur.

In order to submit your possible errata, please email it to our Customer Service Team at wileysupport@wiley.com with the subject line "Possible Book Errata Submission."

Part I Domain 1: Security Principles

[Chapter 1 Confidentiality, Integrity, Availability, and Non-repudiation](#)

[Chapter 2 Authentication, Authorization, and Accounting \(AAA\)](#)

[Chapter 3 Privacy](#)

[Chapter 4 Risk Management](#)

[Chapter 5 Governance Concepts](#)

[Chapter 6 Cybersecurity Controls](#)

[Chapter 7 Professional and Ethical Conduct](#)

Security Principles is the first domain of ISC2's Certified in Cybersecurity exam. It provides the foundational knowledge that anyone in information technology needs to understand as they begin their careers. The domain includes the following five objectives:

1.1 Understand Cybersecurity Concepts

1.2 Understand Risk Management Concepts

1.3 Understand Governance Concepts

1.4 Understand Cybersecurity Controls

1.5 Maintain Professional and Ethical Conduct

Questions from this domain make up approximately 24% of the questions on the CC exam.

CCSM Certified in Cybersecurity Study Guide, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 1 Confidentiality, Integrity, Availability, and Non-repudiation: Objective 1.1

Understand Cybersecurity Concepts

Information plays a vital role in modern business, and we are entrusted with sensitive information about our customers, employees, internal operations, and other critical matters. As cybersecurity professionals, we must work with information security teams, other technology professionals, and business leaders to protect the information we manage.

In this chapter, you'll learn about four of the subobjectives of CC objective 1.1. The remaining material for this objective is covered in [Chapter 2, “Authentication, Authorization, and Accounting \(AAA\)”](#) and [Chapter 3, “Privacy.”](#) The following subobjectives are covered in this chapter:

Confidentiality

Integrity

Availability

Non-repudiation

The CIA Triad

Information security professionals have three key goals when it comes to protecting information and systems. They want to ensure that private data remains secret (confidentiality), that information isn't altered or destroyed without permission (integrity), and that information is available to authorized users when they need it (availability). You can remember these three key goals by thinking of the CIA triad, as shown in [Figure 1.1](#). Each side of this triangle covers one of the three key goals.



Figure 1.1 The CIA triad summarizes the three key goals of information security: confidentiality, integrity, and availability.

Confidentiality Threats

Confidentiality ensures that only authorized individuals have access to information and resources. This is what most people think of when they think about information security—keeping secrets away from prying eyes. And it is, in fact, how information security professionals spend the majority of their time.

As you prepare for the exam, you'll need to understand the main threats against each of the information security goals. I'll talk about many different kinds of threats in this book, but I'll begin with the following: social engineering, phishing, shoulder surfing, snooping, device theft, dumpster diving, eavesdropping, wiretapping, and pretexting.

Social Engineering

Social engineering is a technique where attackers manipulate people into divulging sensitive information or performing actions that compromise security. Rather than exploiting a technical vulnerability, social engineers exploit human nature, using tactics such as impersonation, pretexting, and building false trust to trick their targets. Organizations should train employees to recognize social engineering attempts and verify the identity of anyone making unusual requests.

Phishing

Phishing attacks use fraudulent emails or messages that appear to come from a trusted source, tricking the recipient into clicking a malicious link, opening an infected attachment, or providing sensitive information. Variations include spear phishing, which targets a specific individual (or group of individuals), and whaling, which targets senior executives. Users should carefully inspect sender addresses and avoid clicking links in unexpected messages.

Shoulder Surfing

Shoulder surfing occurs when an attacker looks over someone's shoulder to observe sensitive information, such as passwords, PINs, or confidential documents displayed on a screen. This can happen in offices, airports, coffee shops, or any public place. Privacy screens, which limit the viewing angle of a display, are an effective countermeasure against shoulder surfing.

Snooping

Snooping is exactly what the name implies. The individual engaging in snooping wanders around your office or other facility and simply looks for information they can gather. When people leave sensitive papers on their desks or in a public area, it creates an opportunity for snooping.

Organizations can protect against snooping by enforcing a clean desk policy. Employees should maintain a clean workspace and put away any sensitive materials whenever they step away, even for a moment.

Device Theft

Device theft is a straightforward but serious threat: an attacker physically steals a laptop, smartphone, tablet, or other device to gain access to the data stored on it. Stolen devices may contain sensitive files, cached credentials, or access tokens that allow the thief to breach organizational systems. Full disk encryption, strong authentication, and remote wipe capabilities are essential defenses against device theft.

Dumpster Diving

Dumpster diving attacks also look for sensitive materials, but the attacker doesn't walk around the office; instead, they look through the trash, trying to find sensitive documents that an employee threw in the garbage or recycling bin.

You can protect your organization against dumpster diving attacks with a simple piece of technology: a paper shredder! If you shred documents before discarding them, you'll protect them from dumpster divers.

Eavesdropping

Eavesdropping attacks come in both physical and electronic types. In a physical eavesdropping attack, the attacker simply positions themselves where they can overhear conversations, such as in a cafeteria or hallway, and then listens for sensitive information.

You can protect against eavesdropping by putting rules in place that limit where sensitive conversations may take place. For example, sensitive conversations should generally take place in a closed office or conference room.

Wiretapping

Electronic eavesdropping attacks are also known as *wiretapping*. They occur when an attacker gains access to a network and monitors the electronic data sent within an office.

The best way to protect against electronic eavesdropping is to use encryption to secure information sent over the network. If data is encrypted, an attacker who intercepts it won't be able to make sense of it. I'll talk more about how encryption works later in this book.

Pretexting

The last type of confidentiality attack I'll talk about is *pretexting*. In a pretexting attack, the attacker might pretend that they're on an urgent assignment from a senior leader or impersonate an IT technician.

It's difficult to protect against pretexting attacks. The best defense against these attacks is educating users to recognize the dangers of pretexting and empowering them to intervene when they suspect an attack is taking place.

Integrity Threats

Information security professionals are also responsible for protecting the *integrity* of an organization's information. This means that there aren't any unauthorized changes to or destruction of information. Unauthorized changes may come in the form of a hacker seeking to intentionally alter or destroy information or a service disruption accidentally affecting data stored in a system. In either case, it's the information security professional's responsibility to prevent these lapses in integrity.

This section covers four types of integrity attacks: data modification, impersonation, man-in-the-middle (MitM), and replay.

Data Modification Attacks

Data modification attacks occur when someone gains access to a system and makes changes that violate a security policy. This might be an external attack, such as an intruder breaking into a financial system and issuing themselves checks, or an internal attack, such as an employee increasing their own salary in the payroll system.

Following the principle of *least privilege* is the best way to protect against integrity attacks. Organizations should carefully consider the permissions each employee needs to perform their job and then limit them to the smallest set possible.

Impersonation Attacks

In an *impersonation* attack, the attacker pretends to be someone other than who they actually are. They might impersonate a manager, executive, or IT technician in order to convince someone to change data in a system. The best defense against this type of attack is strong user education.

Man-in-the-middle Attacks

Sometimes impersonation attacks are electronic. In a *man-in-the-middle* (*MitM*) *attack*, the attacker intercepts network traffic as a user logs in to a system and impersonates the system to the user while acting as the user to the system. They then sit in the middle of the communication, relaying information between the user and the system while they monitor everything that is occurring. In this type of attack, the attacker might be able to steal a user's password and use it later to log in to the system themselves.

Replay Attacks

In a *replay attack*, the attacker doesn't get in the middle of the communication; instead, they observe a legitimate user logging in to a system. They then capture the information used to log in to the system and later replay it on the network to gain access themselves.

The best defense against both replay and MitM attacks is encryption to protect communications. For example, web traffic might use the Transport Layer Security (TLS) protocol to prevent eavesdropping on network traffic. You'll learn more about this technology in [Chapter 18, "Encryption."](#)

Availability Threats

As a security professional, you must also understand how to apply security controls that protect the *availability* of information and systems. As the third leg of the CIA triad, availability controls ensure that information and systems remain available to authorized users when needed. They protect against disruptions to normal system operation or data availability.

This section covers five types of events that can disrupt system availability: denial-of-service attacks, power outages, hardware failures, equipment destruction, and service outages.

Denial-of-service Attacks

Denial-of-service (*DoS*) *attacks* occur when a malicious individual bombards a system with an overwhelming amount of network traffic. The idea is to

simply send so many requests to a server that it is unable to answer any requests from legitimate users. In a *distributed denial-of-service (DDoS) attack*, the attacker uses many systems to send overwhelming traffic, making the attack even harder to stop.

You can protect your systems against DoS attacks by using firewalls that block illegitimate requests and by partnering with your Internet service provider to block DDoS attacks before they reach your network.

Power Outages

Power outages can occur at the local or regional level for many reasons. Increased demand can overwhelm the power grid; natural disasters can disrupt service; and other factors can cause power outages, disrupting access to systems.

You can protect against power outages by installing redundant power sources and backup generators that supply power to your system when commercial power is unavailable.

Hardware Failures

Hardware failures can and do occur. Servers, hard drives, network gear, and other equipment all fail occasionally, disrupting access to information. That's an availability problem.

You can protect against hardware failures by building a system that has built-in redundancy so that if one component fails, another is ready to pick up the slack.

Equipment Destruction

Sometimes, equipment is just outright destroyed. This might be the result of intentional or accidental physical damage, or of a larger disaster, such as a fire or a hurricane.

You can protect against small-scale destruction with redundant systems. If you want to protect against larger-scale disasters, you may need to have backup

data centers in remote locations or in the cloud that can keep running when your primary data center is disrupted.

Service Outages

Finally, sometimes service outages occur. This might be due to programming errors, the failure of underlying equipment, or many other reasons. These outages disrupt user access to systems and information and, therefore, constitute an availability concern.

You can protect against service outages by building systems that are resilient to errors and hardware failures.

Non-repudiation

Another important focus of some security controls is providing *non-repudiation*. Repudiation is the act of denying that something is true. Non-repudiation is a security goal that prevents someone from falsely denying that something is true.

For example, you might agree to pay \$10,000 to someone in exchange for a car. If you had only a handshake agreement, you might later be able to repudiate your actions. You might claim that you never agreed to purchase the car or that you agreed to pay a lower price.

To solve this issue, a signed contract is used when a car is sold. Your signature on the document is proof that you agreed to the terms, and if you later go to court, the person selling you the car can prove it by showing the judge the signed document. Physical signatures provide non-repudiation on contracts, receipts, and other paper documents.

There's also an electronic form of the physical signature. *Digital signatures* use encryption technology to provide non-repudiation for electronic documents. You'll learn more about that in [Chapter 18](#).

There are other ways that you can provide non-repudiation as well. You might use biometric security controls, such as a fingerprint or facial recognition, to prove that someone was in a facility or performed an action. You might also use video surveillance for that same purpose. All of these controls enable you

to prove that someone was in a particular location or performed an action, offering some degree of non-repudiation.

Exam Essentials

The CIA triad references the three key goals of information security: confidentiality, integrity, and availability.

Confidentiality protects sensitive information from unauthorized access. The major threats to confidentiality include social engineering, phishing, shoulder surfing, snooping, device theft, dumpster diving, eavesdropping, wiretapping, and pretexting.

Integrity protects information and systems from unauthorized modification or destruction. The major threats to integrity include data modification, impersonation, man-in-the-middle, and replay attacks.

Availability ensures that authorized users have access to information when they need it. The major threats to availability include denial-of-service attacks, power outages, hardware failures, equipment destruction, and service outages.

Non-repudiation uses technical measures to ensure that a user cannot later deny taking an action.

Practice Question 1

Which one of the following security threats would most likely be considered an availability issue?

- A. Replay attack
- B. Power outage
- C. Social engineering
- D. Snooping

Practice Question 2

What are the three key goals of information security programs?

- A. Confidentiality, integrity, and availability
 - B. Confidentiality, integrity, and authorization
 - C. Confidentiality, infrastructure, and authorization
 - D. Communications, infrastructure, and authorization
-

Practice Question 1 Explanation

Availability issues affect authorized users' ability to access the information, systems, or other resources they need. All of the issues listed here are cybersecurity threats that you need to be aware of when you take the CC exam, but only power outages are classified as an availability threat. This is because a power outage can easily disrupt access to systems and information by bringing them offline.

Replay attacks allow an unauthorized individual to impersonate a legitimate user and are primarily considered integrity threats.

Social engineering and snooping attacks may allow an attacker to gain access to sensitive information and are primarily considered confidentiality threats.

Correct Answer: B. Power outage

Practice Question 2 Explanation

The three key goals of any information security program are protecting the confidentiality, integrity, and availability of systems and information.

Confidentiality ensures that only authorized individuals have access to information and resources. Integrity protects information from unauthorized changes. Availability ensures that information and systems are available for authorized use.

Correct Answer: A. Confidentiality, integrity, and availability

CCSM Certified in Cybersecurity Study Guide, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 2 Authentication, Authorization, and Accounting (AAA): Objective 1.1 Understand Cybersecurity Concepts

As a cybersecurity professional, one of the most important things you do is to ensure that only authorized individuals gain access to information, systems, and networks under your protection. You use access controls to provide this assurance.

In this chapter, you'll learn about the subobjectives of CC objective 1.1 that deal with authentication and access control. The remaining material for this objective is covered in [Chapter 1, “Confidentiality, Integrity, Availability, and Non-repudiation,”](#) and [Chapter 3, “Privacy.”](#) The following subobjective is covered in this chapter:

Authentication, Authorization, Accounting (AAA)

Access Control Process

The access control process consists of three steps that you must understand as you prepare for the Certified in Cybersecurity (CC) exam: identification, authentication, and authorization. Access control systems also perform another important task: accounting.

Exam Tip

Together, the functions of an access control system are referred to as AAA, or “triple-A.” Those three As are for authentication, authorization, and accounting. Yes, that leaves identification off the list, but identification is assumed to be part of the process if you're performing authentication, and adding an “I” would ruin the easy acronym!

Identification

During the first step of the process, *identification*, an individual makes a claim about their identity. The person trying to gain access doesn't present any proof at this point; they simply make an assertion. It's important to remember that the identification step is only a claim, and the user could certainly be making a false claim! Imagine a physical-world scenario in which you want to enter a secure office building for an appointment. During the identification step of the process, you might approach the security desk and say, "Hi, I'm Mike Chapple."

Authentication

Proof comes into play during the second step of the process: *authentication*. During the authentication step, the individual proves their identity to the access control system. In the office building example, the guard would likely want to see my driver's license to confirm my identity.

Authorization

Just proving your identity isn't enough to gain access to a system, however. The access control system also needs to verify that you are *allowed* to access it. That's the third step of the access control process: *authorization*. In the office building example, the security guard might check a list of that day's appointments to see if my name is on it.

Accounting

In addition to identification, authentication, and authorization, access control systems also provide an *accounting* functionality that allows administrators to track user activity and reconstruct it from logs. In the office building example, the security guard may record my access to the building in a visitor logbook.

[Figure 2.1](#) shows this process all coming together.



Figure 2.1 The physical access control process.

Digital Access Control

So far, I've talked about identification, authentication, authorization, and accounting in the context of gaining access to a building. Now I'll discuss how they work in the electronic world. When we log in to a system, we most often identify ourselves using a username, typically a combination of letters from our names. When we reach the authentication phase, we're commonly asked to enter a password.

Note

There are many other ways to authenticate, and I'll discuss them later in this chapter. I'll also discuss how strong access control systems combine multiple authentication approaches.

In the electronic world, authorization often takes the form of *access control lists* that itemize the specific file system permissions granted to an individual user or a group of users. Users go through identification, authentication, and authorization when requesting access to a resource.

Digital systems also perform a significant amount of accounting. This may include tracking user activity on systems and even logging user web browsing history. Any tracking conducted as part of an organization's monitoring program should comply with the law and the organization's privacy policy.

[Figure 2.2](#) shows the digital access control process coming together.



Figure 2.2 The digital access control process.

Password Policies

When you set a password policy for your organization, you have several technical controls that let you specify requirements for how users choose and maintain their passwords. This section discusses a few of those mechanisms.

Password Length

The simplest and most common control on passwords is setting the *password length*. This is the minimum number of characters required in a password. Many organizations require passwords to be at least 8 characters long, though current NIST guidelines recommend a minimum of 15 characters for stronger security. The longer a password is, the harder it is to guess.

Password Complexity

Organizations may also set *password complexity* requirements. These requirements force users to include various character types in their passwords, such as uppercase and lowercase letters, digits, and special characters. Just as with password length, the more character types a password contains, the harder it is to guess. However, current NIST guidelines recommend against imposing complexity requirements, as they often lead users to choose weaker, more predictable patterns. Instead, NIST recommends checking passwords against lists of commonly used or compromised passwords.

Password Expiration

Password expiration requirements force users to change their passwords periodically. For example, an organization might set a password expiration period of 180 days, forcing users to change their passwords every 6 months. These days, many organizations no longer have password expiration requirements, allowing users to keep the same password for as long as they'd like and requiring them to change it only if it's compromised.

Password History

Password history requirements are designed to prevent users from reusing their previous passwords. Organizations with password history requirements configure their systems to remember each user's previous passwords and prevent reuse. Password history controls allow the administrator to identify how many previously used passwords are remembered for each user.

Password Changes

Every organization should allow users to change their passwords quickly and easily. You want users to be able to privately select their own passwords whenever they are concerned that their passwords may be compromised.

Password Reset

A *password reset* is a recovery method that enables users to create a new password after forgetting their password. Organizations should carefully evaluate their password reset process for users who forget their passwords. If they're not well designed, these processes can give attackers an opportunity to gain unauthorized access to a system by performing password resets.

Password Reuse

IT teams should also strongly encourage users not to reuse passwords across multiple sites. This is difficult to actually enforce, but it does provide a strong measure of security. If a user reuses the same password across many sites and one of those sites is compromised, an attacker might test that password on other sites, hoping the password owner reuses it.

Password Managers

It's difficult for users to manage unique passwords for every site they visit. That's where password managers play a crucial role. These valuable tools are secure password vaults, often protected by biometric security mechanisms that create and store unique passwords. They then automatically fill those passwords on websites when the user visits them. That way, users can have unique, strong passwords for every site they visit without having to remember them all.

[Figure 2.3](#) shows an example of LastPass, a popular password manager, being used to create a new, strong password.

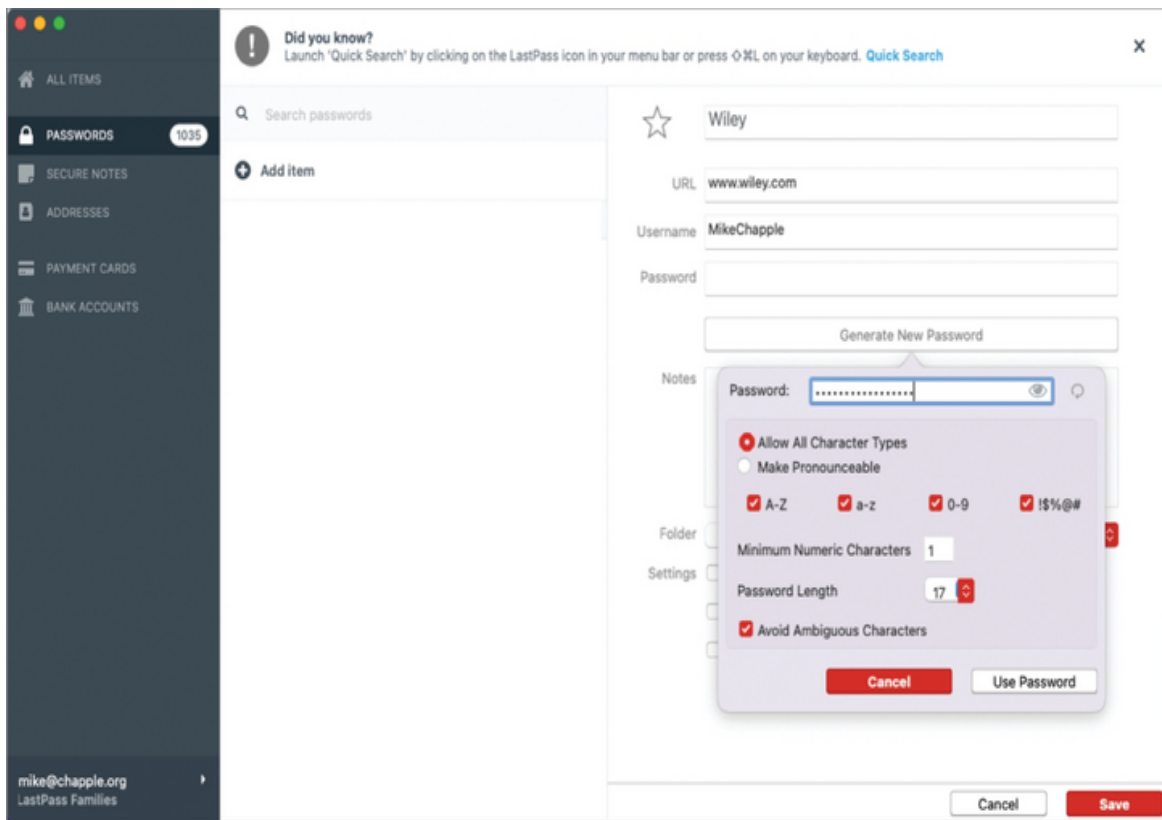


Figure 2.3 Creating a password in LastPass.

Authentication Factors

Computer systems offer many different authentication techniques that allow users to prove their identity. This section looks at the following three authentication factors:

Something you know

Something you are

Something you have


Something You Know

By far, the most common authentication factor is *something you know*. I've already talked about how passwords are the most common authentication method. This is a "something you know" factor because the password is

something that the user remembers and enters into a system during the authentication process. Personal identification numbers (PINs) and the answers to security questions are also examples of something you know.

Something You Are

The second authentication factor is *something you are*. *Biometric* authentication methods measure one of your physical characteristics, such as a fingerprint, eye pattern, face, or voice. [Figure 2.4](#) shows an example of a smartphone fingerprint reader performing biometric authentication.



[Figure 2.4](#) Fingerprint authentication on a smartphone.

Source: artimedvedev/Adobe Stock Photos

[Figure 2.5](#) shows a more complex method—an iris scan—being performed as a biometric access control for entering a facility.



Figure 2.5 Iris scan authentication for entering a facility.

Source: demphoto/Adobe Stock Photos

Something You Have

The third authentication factor, *something you have*, requires the user to have physical possession of a device, such as a smartphone running an authenticator app or a hardware token. These devices generate one-time passcodes that are displayed to the user and allow them to prove that they have access to a physical device. Similarly, smart cards can serve as a “something you have” factor, requiring that the user insert a card with a digital chip into a specialized reader.

Multi-factor Authentication

When used alone, any one of these methods provides some security for systems. However, each has its own drawback. For example, an attacker might steal a user's password through a phishing attack. Once they have the password, they can use it to assume the user's identity. Other authentication factors aren't foolproof either. If you use smart card authentication to implement something you have, the user may lose the smart card. Someone coming across it may then impersonate the user.

The solution to this problem is to combine two or more authentication factors, such as something you know and something you have. This approach is known as *multi-factor authentication (MFA)*.

Take the two methods just discussed: passwords and smart cards. When used alone, either one is subject to attackers gaining knowledge of the password or stealing a smart card. However, if an authentication system requires both a password (something you know) and a smart card (something you have), it brings added security. If the hacker steals the password and doesn't have the required smart card, or vice versa, access is prevented. It suddenly becomes much more difficult for the attacker to gain access to the account. Because something you know and something you have are different factors, this is MFA.

You can also combine other factors. For example, a fingerprint reader (something you are) might also require a PIN (something you know). That's another example of MFA.

When evaluating MFA, it's important to remember that the factors must be *distinct*. An approach that combines a password with the answer to a security question is *not* MFA, because both factors are something you know.

Exam Tip

When you take the CC exam, you'll likely find a question about MFA. Be sure the authentication techniques use two different factors. Mistaking two "something you know" factors for MFA is a common exam mistake!

Exam Essentials

The access control process consists of three major steps. Identification is when a user asserts their identity. Authentication is the process by which a user proves their identity. Authorization is the process by which the system determines whether a user is allowed to perform a requested action.

Accounting processes create a record of who performed which actions on a system and are useful when investigating security incidents.

Multi-factor authentication (MFA) combines at least two of the three factors: something you know, something you have, and something you are.

Password length requirements set a minimum number of characters a user's password must contain, whereas password complexity requirements mandate the use of different character types.

Password history requirements prevent the reuse of old passwords, whereas password expiration requirements require periodic password changes. Users should be permitted to reset their passwords whenever they wish.

Users should be discouraged from reusing passwords across multiple sites, as this increases the risk of compromise. Password managers provide a convenient way to manage many unique, strong passwords.

Practice Question 1

You are considering deploying a multi-factor authentication system to protect access to your organization's virtual private network (VPN). Which one of the following combinations of access controls would meet this requirement?

- A. Password and PIN
- B. Fingerprint and iris scan
- C. Smart card and fingerprint
- D. Key fob and smart card

Practice Question 2

Andy is attempting to change his password and has created the following long password:

p7djknqr2LAD

He receives an error message that his password must contain a symbol. Which password policy is he failing to meet?

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

Practice Question 1 Explanation

Before you can answer this question, you must identify the authentication factors in use here.

Passwords and PINs are both examples of something you know.

Fingerprints and iris scans are both examples of something you are.

Key fobs and smart cards are both examples of something you have.

The only answer option that combines factors from two different categories is the use of a smart card (something you have) with a fingerprint (something you are).

Correct Answer: C. Smart card and fingerprint

Practice Question 2 Explanation

This is an example of a password complexity requirement. The message is requiring that Andy use a symbol in his password, and password complexity requirements require the use of different character types.

There is no indication that Andy is attempting to reuse an old password, which would violate a password history or reuse requirement. There is also no indication that his password is too short and does not meet the password length requirements.

Correct Answer: C. Password complexity

CCSM Certified in Cybersecurity Study Guide, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 3 Privacy: Objective 1.1 Understand Cybersecurity Concepts

In the digital era, the organizations we deal with collect extensive information about individuals and their actions. From credit card transactions to educational records, each of us generates a significant trail of data, and we are rightfully concerned about the *privacy* of that information.

In this chapter, you'll learn about the privacy subobjective of CC objective 1.1. The remaining material for this objective is covered in [Chapter 1](#), "[Confidentiality, Integrity, Availability, and Non-repudiation](#)," and [Chapter 2](#), "[Authentication, Authorization, and Accounting \(AAA\)](#)." The following subobjective is covered in this chapter:

Privacy

Privacy

As a cybersecurity professional, you have a few interests in how organizations collect and use personal information:

You are obviously concerned about the privacy of your own personal information.

You have a responsibility to educate users in your organization about how they can protect their own personal information.

You have a responsibility to assist the privacy officials within your organization with the work they need to do to protect the personal information entrusted to your organization.

Types of Private Information

Private information may come in many forms. Two of the most common elements of private information are personally identifiable information and protected health information:

Personally identifiable information (PII) includes any information that can be linked to a specific individual.

Protected health information (PHI) includes health care records that are regulated under the Health Insurance Portability and Accountability Act (HIPAA).

Expectation of Privacy

Privacy programs are based on the legal principle of *reasonable expectation of privacy*. Many laws governing whether information must be protected are based on whether the person disclosing the information had a reasonable expectation of privacy when they did so, and whether the disclosure would violate that expectation.

When you put content on social media, whether it be a post on a social networking site or a comment on a shared video, you generally have no reasonable expectation of privacy. You're posting that content publicly or to a large group, and a reasonable person would assume it is not a private conversation.

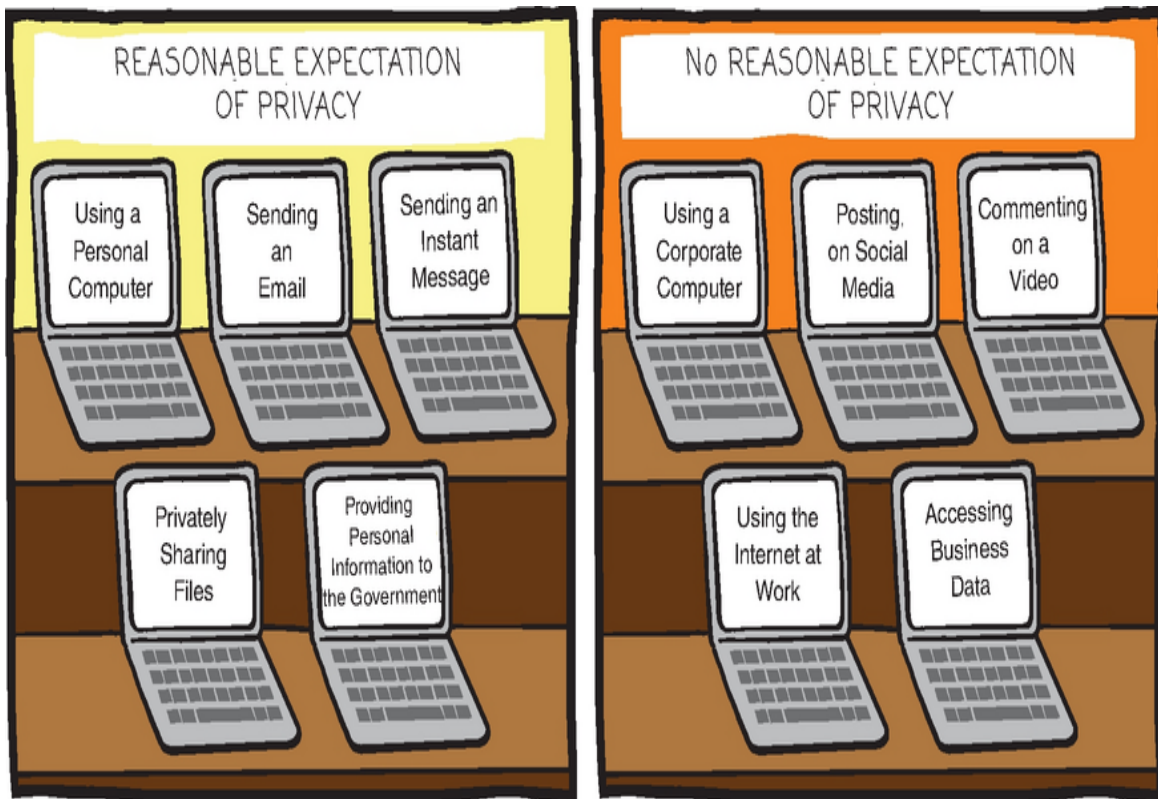
However, when you send a message to someone over email or instant messaging, or share a file through a file-sharing site with a specific person or small group, you do have an expectation of privacy. If you don't encrypt the message, you run the risk that others will eavesdrop, but you have a more reasonable expectation of privacy.

At the other end of the spectrum, when you provide private information to a government agency through their website, such as completing your taxes or registering for a medical insurance program, you have a much greater expectation of privacy. In those cases, you expect that the agency collecting information will treat that information with care and not share it without your permission. And, in fact, most countries have laws that require this information to be kept private.

When you're using a computer or network that belongs to your employer, you generally do not have a reasonable expectation of privacy. The employer owns that equipment and is normally legally entitled to monitor its use. If you're using software on your employer's desktop computer, accessing business information over a corporate network, or even using the Internet at work for

personal use, you generally should not expect to have privacy for your communication.

As a cybersecurity professional, you should clearly and accurately communicate users' privacy expectations to them. It's important to reinforce that when employees of your organization use systems or networks owned by the organization, they should not have an expectation of privacy. [Figure 3.1](#) illustrates cases where users do and do not have a reasonable expectation of privacy.



[Figure 3.1](#) Do you have a reasonable expectation of privacy?

You also need to ensure that you and your organization handle confidential information appropriately. In addition to protecting personal information about employees and customers, you should also safeguard other sensitive information, including passwords and the company's confidential information.

Privacy Management Framework

The *Privacy Management Framework (PMF)* aims to establish a global framework for privacy management. The PMF includes nine components developed by the American Institute of Certified Public Accountants (AICPA) with subject-matter expert input.

The nine PMF components are as follows:

- . Management
- . Agreement, notice, and communication
- . Collection and creation
- . Use, retention, and disposal
- . Access
- . Disclosure to third parties
- . Security for privacy
- . Data integrity and quality
- . Monitoring and enforcement

The remainder of this section explores each of these components in more detail.

Management

Management is the first of the nine privacy components, and the PMF defines it as follows: “The entity defines, formally documents, communicates, and assigns responsibility and accountability for its PI (personal information) privacy policies and procedures.” The criteria that organizations should follow to establish control over the management of their privacy program include the following:

Creating written privacy policies and communicating those policies to personnel

Assigning responsibility and accountability for those policies to a person or team

Establishing procedures for the review and approval of privacy policies and changes to those policies

Ensuring that privacy policies are consistent with applicable laws and regulations

Performing privacy risk assessments on at least an annual basis

Ensuring that contractual obligations to customers, vendors, and partners are consistent with privacy policies

Assessing privacy risks when implementing or changing technology infrastructure

Creating and maintaining a privacy incident management process

Conducting privacy awareness and training and establishing qualifications for employees with privacy responsibilities

Agreement, Notice, and Communication

The second component, *agreement, notice, and communication*, requires that organizations inform individuals about their privacy practices and obtain agreement for information processing. The PMF defines this component as follows: “The entity makes formal agreements, notifies and communicates with and offers choices when seeking data subject consents, including reasons why and purposes for which the entity seeks to obtain and use a data subject’s PI.”

The component incorporates the following criteria:

Including notice practices in the organization’s privacy policies

Notifying individuals about the purpose of collecting personal information and the organization’s policies surrounding the other components

Providing notice to individuals at the time of data collection, when policies and procedures change, and when the organization intends to use information for new purposes not disclosed in earlier notices

Writing privacy notices in plain and simple language and posting them conspicuously

Including choice and consent practices in the organization’s privacy policies

Informing individuals about the choice and consent options available to them and the consequences of refusing to provide personal information or

withdrawing consent to use personal information

Obtaining implicit or explicit consent at or before the time that personal information is collected

Notifying individuals of proposed new uses for previously collected information and obtaining additional consent for those new uses

Obtaining direct explicit consent from individuals when the organization collects, uses, or discloses sensitive personal information

Obtaining consent before transferring personal information to or from an individual's computer or device

Collection and Creation

The component of *collection and creation* governs how organizations come into possession of personal information. The PMF defines this component as follows: "The entity collects and creates PI only for the purposes identified in its agreements with data subjects, and in ongoing communications with and notices provided to data subjects."

The criteria associated with the collection and creation component are as follows:

Including collection practices in the organization's privacy policies

Informing individuals that their personal information will be collected only for identified purposes

Including details on the methods used to collect data and the types of data collected in the organization's privacy notice

Collecting information using fair and lawful means and only for the purposes identified in the privacy notice

Confirming that any third parties who provide the organization with personal information have collected it fairly and lawfully and that the information is reliable

Informing individuals if the organization obtains additional information about them

Note

Although it is not explicitly included in the collection criteria, data minimization is another crucial component of privacy programs. This principle says that an organization should collect the minimum amount of personal information necessary to meet their objectives and discard that information when it is no longer needed for that purpose.

Use, Retention, and Disposal

Organizations must maintain the privacy of personal information throughout its life cycle. This is where the component of *use, retention, and disposal* plays an important role. The PMF defines this component as follows: “The entity limits the use of PI to the purposes identified in the formal agreements/notices, and for which a data subject has provided explicit (or implicit) consent. The entity retains PI for the time necessary to fulfill the stated purposes identified in the formal agreements/notices, or as required by law or regulation. Once those purposes have been met, the entity securely disposes of the information.”

The criteria associated with the use, retention, and disposal component is as follows:

Including collection practices in the organization’s privacy policies

Informing individuals that their personal information will be used only for disclosed purposes for which the organization has obtained consent, and then abiding by that statement

Informing individuals that their data will be retained for no longer than necessary, and then abiding by that statement

Informing individuals that information that is no longer needed will be disposed of securely, and then abiding by that statement

Access

One of the core elements of individual privacy is the belief that individuals should have the right to access the information an organization holds about them and, when necessary, to correct it. This right to correct information is also known as the right of redress. The PMF definition of the *access* component is as follows: “The entity provides data subjects with access to their PI when requested or when asked to update and correct data errors or make changes.”

The criteria associated with the access component are as follows:

Including practices around access to personal information in the organization’s privacy policies

Informing individuals about the procedures for reviewing, updating, and correcting their personal information

Providing individuals with a mechanism to determine whether the organization maintains personal information about them and to review any such information

Authenticating an individual’s identity before providing them with access to personal information

Providing access to information in an understandable format within a reasonable period of time, either for a reasonable charge that is based on the organization’s actual costs or at no cost

Informing individuals in writing why any requests to access or update personal information were denied, and informing them of any appeal rights they may have

Providing a mechanism for individuals to update or correct personal information and providing that updated information to third parties who received it from the organization

Disclosure to Third Parties

Some challenging privacy issues arise when organizations maintain personal information about an individual and then choose to share it with third parties in the course of business. The PMF defines the *disclosure to third parties* component as follows: “The entity discloses PI to third parties only for the

purposes identified in data subject privacy agreements and its privacy notice and with the explicit consent of the data subject.”

The criteria associated with the disclosure to third parties component are as follows:

Including third-party disclosure practices in the organization’s privacy policies

Informing individuals of any third-party disclosures that take place and the purpose of those disclosures

Informing third parties who receive personal information from the organization that they must comply with the organization’s privacy policy and handling practices

Disclosing personal information to third parties without notice or for purposes other than those disclosed in the notice, only when required to do so by law

Disclosing information to third parties only under the auspices of an agreement that the third party will protect the information consistent with the organization’s privacy policy

Implementing procedures designed to verify that the privacy controls of third parties receiving personal information from the organization are functioning effectively

Taking remedial action when the organization learns that a third party has mishandled personal information shared by the organization

Security for Privacy

Protecting the security of personal information is deeply intertwined with protecting its privacy. Organizations can’t provide individuals with assurances about how they handle personal data if they can’t protect that information from unauthorized access. The PMF defines *security for privacy* as follows: “The entity protects PI against unauthorized access, removal, alteration, destruction and disclosure (both physical and logical).”

The criteria associated with the security for privacy component are as follows:

Including security practices in the organization’s privacy policies

Informing individuals that the organization takes precautions to protect the privacy of their personal information

Developing, documenting, and implementing an information security program that addresses the major privacy-related areas of security

Restricting logical access to personal information through the use of strong identification, authentication, and authorization practices

Restricting physical access to personal information through the use of physical security controls

Protecting personal information from accidental disclosure due to natural disasters and other environmental hazards

Applying strong encryption to any personal information that is transmitted over public networks

Avoiding the storage of personal information on portable media, unless absolutely necessary, and using encryption to protect any personal information that must be stored on portable media

Conducting periodic tests of security safeguards used to protect personal information

Data Integrity and Quality

When we think about protecting the privacy of personal information, we often first consider the proper collection and use of that information, as well as potential unauthorized disclosure. However, it's also important to consider the accuracy of that information. Individuals may be harmed by incorrect information just as much, if not more, as by improperly handled information. The PMF *data integrity and quality* component states that "The entity maintains accurate, complete and relevant PI for the purposes identified in the notice and protects the representational integrity of the PI in its ongoing interactions with data subjects."

The criteria associated with the data integrity and quality component are as follows:

Including data quality practices in the organization's privacy policies

Informing individuals that they bear responsibility for providing the organization with accurate and complete personal information and informing the organization if corrections are required

Maintaining personal information that is accurate, complete, and relevant for the purposes for which it will be used

Monitoring and Enforcement

Privacy programs are not a one-time project. It's true that organizations may make a substantial initial investment of time and energy to build up their privacy practices, but those practices must be monitored over time to ensure that they continue to operate effectively and meet the organization's privacy obligations as business needs and information practices evolve. The PMF *monitoring and enforcement* component states that "The entity monitors compliance with its privacy policies and procedures and has procedures to address privacy-related complaints and disputes."

The criteria associated with the monitoring and enforcement component are as follows:

Including monitoring and enforcement practices in the organization's privacy policies

Informing individuals about how they should contact the organization if they have questions, complaints, or disputes regarding privacy practices

Maintaining a dispute resolution process that ensures that every complaint is addressed and that the individual who raised the complaint is provided with a documented response

Reviewing compliance with privacy policies, procedures, laws, regulations, and contractual obligations on an annual basis

Developing and implementing remediation plans for any issues identified during privacy compliance reviews

Documenting cases where privacy policies were violated and taking any necessary corrective action

Performing ongoing monitoring of the privacy program based on a risk assessment

Exam Essentials

IT professionals have an obligation to protect the privacy of personal information, particularly when the individuals involved have a reasonable expectation of privacy.

The types of personal information that must be safeguarded include personally identifiable information (PII), which is any personal information that may be tied to an individual, and protected health information (PHI), which includes medical records.

The nine Privacy Management Framework (PMF) components are as follows:

1. Management
2. Agreement, notice, and communication
3. Collection and creation
4. Use, retention, and disposal
5. Access
6. Disclosure to third parties
7. Security for privacy
8. Data integrity and quality
9. Monitoring and enforcement

Practice Question 1

Which of these people would have the least expectation of privacy?

- A. A patient discussing medical records with a nurse
 - B. An employee sending an email to a friend over the office network
 - C. A student discussing grades with a teacher
 - D. Two people having a private conversation
-

Practice Question 2

Norm is reviewing his organization's privacy practices and observes that the privacy notice is not posted on their website in a location that is accessible to customers. Which PMF component is most directly violated by this action?

- A. Agreement, notice, and communication
 - B. Choice and consent
 - C. Communication
 - D. Collection and creation
-

Practice Question 1 Explanation

A patient discussing medical records with a nurse is in a situation where there is a strong expectation of privacy. These records are protected health information (PHI) and should be treated very carefully. Of the situations listed, this is likely the one with the greatest expectation of privacy.

A student discussing grades with a teacher and two people having a private conversation are also in situations where they should have some expectation of privacy.

An employee of an organization sending an email at work has no expectation of privacy. The employer is entitled to inspect any information sent or received using the office network, even if that information is of a personal nature.

Correct Answer: B. An employee sending an email to a friend over the office network

Practice Question 2 Explanation

The second PMF component—agreement, notice, and communication—requires that organizations inform individuals about their privacy practices. One of the criteria for this component includes writing privacy notices in plain and simple language and posting them conspicuously. Norm's organization is not doing this.

This does not directly impact the component of choice and consent or collection and creation.

Communication seems like an obvious answer here, but it is not one of the nine PMF components.

Correct Answer: A. Agreement, notice, and communication

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 4 Risk Management: Objective 1.2

Understand Risk Management Concepts

Risks abound in the world of information security. From hackers and malware to lost devices and missing security patches, information security professionals have a lot on their plates. Of course, addressing each risk takes both time and money. Risk management is the practice of identifying, assessing, treating, and monitoring the risks facing an organization.

In this chapter, you'll learn about CC objective 1.2. The following subobjectives are covered in this chapter:

Risk management lifecycle

Risk management processes

Organizations manage risk through a structured process known as the *risk management lifecycle*, shown in [Figure 4.1](#). This lifecycle consists of several key phases: first, identifying the risks facing the organization; then, assessing those risks to understand their likelihood and potential impact; next, treating the identified risks by selecting and, finally, implementing appropriate responses. The sections that follow in this chapter walk through each of these phases in detail.



Figure 4.1 The risk management lifecycle.

Risk Categories and Types

The foundation of a cybersecurity professional's work is risk management. Organizations face many kinds of risk, and it's your job to identify, assess, treat, and monitor those risks to protect your information and assets. This section discusses the different kinds of risk that exist in our everyday world.

Internal and External Risks

Risks can be divided into two categories: internal and external. *Internal risks* are those that arise from within the organization. For example, if the way you process checks creates an opportunity for employees in the accounting department to commit fraud, that's an internal risk.

You can often address internal risks by adding *internal controls*. These security measures reduce the likelihood and/or impact of a risk occurring. In the accounting example, adding a two-person control to the issuance of checks might reduce the risk of fraud by requiring two people to cooperate to issue a check.

External risks are those that originate outside the organization. For example, the risk of an attacker targeting your organization with a ransomware attack is external. You can't do much to stop the attacker from attempting the attack, but you can build controls that reduce the likelihood of success and/or the impact if it does occur, such as implementing multi-factor authentication (MFA) or launching a social engineering threat awareness campaign.

Multiparty Risks

The next type of risk is *multiparty risks*. These are risks shared among many organizations. For example, if a software as a service (SaaS) provider is compromised, that is a multiparty risk because it poses a risk to all the customers of that service provider.

Specific Risks

In addition to the general categories of internal, external, and multiparty risks, there are a few specific risk types that you should be familiar with as you prepare for the CC exam. These are the risks associated with legacy systems, intellectual property, and software license compliance.

Legacy Systems

Legacy systems pose a unique type of risk to the organization. These systems have been around for a long time and remain in use despite their age and possible lack of maintenance. It's often difficult to secure older systems, especially those no longer supported by the manufacturer. Any organization using legacy systems should either replace them with a modern solution or carefully design a set of security controls to mitigate the associated risks.

Intellectual Property

In the information age, the value delivered by many businesses resides in their *intellectual property*. If attackers were able to alter, destroy, or steal this information, it would cause significant harm and damage to the business.

Therefore, intellectual property theft poses a risk to information-based organizations.

Software License Compliance

Finally, make sure you consider the risks associated with *software license compliance*. Businesses often go to great lengths to protect their investment in software intellectual property, including conducting audits of organizations and imposing significant fines on those who violate license agreements. It's a good idea to use license monitoring software to manage your software license compliance efforts.

Risk Identification and Assessment

The world of cybersecurity is full of risks. As a result, you have to figure out how to allocate your limited time and resources to address the most important risks first. To accomplish this, you need to prioritize your risk list so that you spend your precious resources where they will have the greatest security impact.

Risk identification and assessment is the process of identifying and triaging risks facing an organization based on their likelihood of occurrence and expected impact.

The Language of Risk

First, we need a common language. In everyday life, people often use the terms *threat*, *vulnerability*, and *risk* interchangeably. However, they are actually three different concepts:

Threats are internal or external forces that jeopardize the security of your information and systems. Threats might be naturally occurring, such as hurricanes and wildfires, or human-made, such as hacking and terrorism. You can't normally control what threats are out there. They exist independently of you and your organization.

Vulnerabilities are weaknesses that threats can exploit to undermine the confidentiality, integrity, or availability of your information or systems. These

might include missing patches, promiscuous firewall rules, or other security misconfigurations. You *do* have control over the vulnerabilities in your environment, and security professionals spend much of their time hunting down and remediating vulnerabilities.

Risks occur when your environment contains both a vulnerability and a corresponding threat that could exploit it, as shown in [Figure 4.2](#). For example, if you haven't patched a database server recently and hackers create malware that exploits an unpatched vulnerability, you face a risk. You are vulnerable because you're missing a security control, and there is a threat: the new malware.



[Figure 4.2](#) Risks are the combination of a threat and a corresponding vulnerability.

There is no risk if either the threat or vulnerability factor is missing. For example, if you live in an area far from the coast, it doesn't matter if your building is vulnerable to hurricanes, because there is no threat of a hurricane in your region. Similarly, if you store your backup tapes in a fireproof box, the risk of them being destroyed in a fire is reduced because the container is not vulnerable to fire.

Exam Tip

There is one related term that you should know for the exam. A *threat vector* is the method that an attacker uses to get to a target. This might be a hacker toolkit, social engineering, or physical intrusion.

Ranking Risks

Once you've identified the risks facing your organization, you're likely still left with a somewhat overwhelming list. The next stage in the process ranks those risks by two factors: likelihood and impact.

Risk Likelihood

The *likelihood* of a risk is the probability that it will actually occur. For example, there is a risk of an earthquake in both California and Wisconsin. When you look at the data, however, you find that the probability of an earthquake occurring is far higher in California, where almost 5,000 significant earthquakes occurred over a 25-year period. During that same time, Wisconsin didn't experience a single major earthquake. Therefore, security professionals in California must be hypervigilant about earthquake risk, while those in Wisconsin can probably ignore it.

Risk Impact

The *impact* of a risk is the amount of damage that will occur if the risk materializes. For example, an earthquake might cause devastating damage to a data center, while a rainstorm might cause none.

Risk Assessment Techniques

In risk assessments, two categories of techniques are available for assessing the likelihood and impact of a risk: qualitative and quantitative.

Qualitative techniques use subjective judgments to assess risks, typically categorizing them as low, medium, or high on both the likelihood and impact scales. [Figure 4.3](#) shows an example of a qualitative risk assessment matrix. When considering a specific risk, the assessor first rates the likelihood as low, medium, or high, then rates the impact. The chart then categorizes the overall risk. For example, a risk with a high likelihood and high impact would be categorized as a critical risk, whereas a risk with a medium likelihood and low impact would be categorized as low risk.

		IMPACT		
		LOW	MEDIUM	HIGH
LIKELIHOOD	HIGH	MEDIUM	HIGH	CRITICAL
	MEDIUM	LOW	MEDIUM	HIGH
	LOW	LOW	LOW	MEDIUM

Figure 4.3 Qualitative risk assessment.

Quantitative techniques use objective numeric ratings to assess the likelihood and impact of a risk. When performing a quantitative risk assessment, you'll do some mathematical calculations to determine the precise amount of financial damage you can expect from a given risk in any typical year.

Exam Tip

You don't need to know how to do the math for a quantitative risk assessment on the CC exam, but you'll learn it if you move on to prepare for the CISSP or other advanced cybersecurity certifications.

Risk Treatment Strategies

Once you complete a risk assessment for your organization, you're left with a prioritized list of risks that require your attention. *Risk treatment* is the process of systematically analyzing potential responses to each risk and implementing strategies to appropriately control them.

No matter the risk you're managing, you have four basic options for addressing it. You can:

Avoid the risk

Transfer the risk

Mitigate the risk

Accept the risk

Risk Avoidance

When you practice *risk avoidance*, you change your organization's business practices so that you are no longer in a position where that risk can affect your business.

For example, imagine that you conducted a risk assessment of the flooding risk to your organization's data center. If you choose to pursue a risk-avoidance strategy for that risk, you might relocate your data center to a facility that is not at risk of flooding.

Risk Transference

Risk transference attempts to shift the impact of a risk from your organization to another organization. The most common example of risk transference is an insurance policy. Many organizations are now considering purchasing cybersecurity insurance policies to protect against financial losses from cyberattacks, including data breaches and identity theft. It's important to remember, however, that you can't always transfer a risk completely. For example, you can purchase insurance to cover the financial damage caused by a security breach, but no insurance policy can repair your business's reputation in the eyes of your customers.

In the flood risk example, you might transfer the financial risk of your data center flooding from your organization to an insurance company by purchasing flood insurance.

Risk Mitigation

Risk mitigation takes actions designed to reduce the likelihood and/or impact of a risk.

If you want to mitigate the risk of flooding in your data center, you might engage a flood control specialist to install systems that divert water away from your facility.

Risk Acceptance

In almost every risk assessment, managers find themselves confronted with a very long list of risks and with inadequate resources to avoid, transfer, or mitigate them all. For business reasons, they must accept some of those risks.

Risk acceptance should occur only as part of a thoughtful analysis that determines that the cost of performing another risk management action outweighs the benefit of controlling the risk.

In the flooding scenario, you might conclude that all other risk management options are too costly and decide to continue operations at your current facility as is, and deal with the aftermath of a flood should it occur.

Risk Profile and Tolerance

Every organization must choose the appropriate mix of risk treatment strategies for its own technical and business environment. The combination of risks that affect an organization is known as its *risk profile*, and the organization adopts risk treatment strategies to address those risks.

[Figure 4.4](#) shows the process an organization follows to manage risk. The initial level of risk that exists in an organization before any controls are put in place is the organization's *inherent risk*. Then controls are applied to reduce that risk, but, of course, not every risk can be completely eliminated. The risk that remains after controls reduce the inherent risk is known as *residual risk*.




Figure 4.4 Applying controls reduces the inherent risk down to the residual risk but also introduces control risk.

The controls themselves may introduce new risks. For example, if you install a firewall as a risk management control, it may substantially reduce your risk,

but it also introduces a new risk: the firewall itself may fail. The new risk introduced by adding controls is called *control risk*.

The reality is that the organization will need to accept some ongoing risk in order to continue operations. Business leaders must decide how much risk they choose to accept. This is a process known as determining the organization's *risk tolerance*.

Exam Tip

The goal of risk management is to ensure that the combination of residual and control risks is below the organization's risk tolerance.

Exam Essentials

Internal risks are those that arise from within the organization. External risks are those that originate outside the organization. Multiparty risks are shared among many different organizations.

Threats are internal or external forces that jeopardize the security of your information and systems. Vulnerabilities are weaknesses in assets that threats can exploit to undermine the confidentiality, integrity, or availability of your information or systems. Risks occur when your environment contains both a vulnerability and a threat that could exploit it.

Risks are prioritized based on their likelihood and impact. The likelihood of a risk is the probability that it will actually occur. The impact of a risk is the amount of damage that will occur if the risk materializes.

Qualitative risk assessment techniques use subjective judgments to assess risks, typically categorizing them as low, medium, or high on both the likelihood and impact scales. Quantitative risk assessment techniques use objective numeric ratings to assess the likelihood and impact of a risk.

The four risk treatment options are avoiding, transferring, mitigating, and accepting a risk.

An organization's risk profile begins with its inherent risk. Security professionals then apply controls to reduce the inherent risk, leaving a residual risk. Controls may introduce new control risks.

Practice Question 1

Jen identified a missing patch on a Windows server that might allow an attacker to gain remote control of the system. After consulting with her manager, she applied the patch. From a risk management perspective, what has she done?

- A. Removed the threat
- B. Reduced the threat
- C. Removed the vulnerability
- D. Reduced the vulnerability

Practice Question 2

Grace recently completed a risk assessment of her organization's exposure to data breaches and determined that there is a high risk of loss of sensitive personal information from a type of attack called SQL injection. She is considering a variety of approaches to managing this risk.

Grace's first idea is to add a web application firewall to protect her organization against SQL injection attacks. Which risk management strategy does this approach adopt?

- A. Risk acceptance
- B. Risk avoidance
- C. Risk mitigation
- D. Risk transference

Practice Question 1 Explanation

Jen can reduce the risk associated with the missing patch by taking one of two courses of action. She can either reduce/remove the threat or reduce/remove the vulnerability.

Jen has no control over what an external attacker does, so she cannot reduce or remove the threat.

However, she does have control over the system's patch status. By applying the patch, Jen has removed the vulnerability from her server. This goes beyond actions that would simply reduce the vulnerability, such as adding an additional firewall.

Correct Answer: C. Removed the vulnerability

Practice Question 2 Explanation

Installing new controls or upgrading existing controls is an effort to reduce the likelihood and/or impact of a risk. This is an example of a risk mitigation activity because it is designed to reduce the likelihood and/or impact of a risk.

If Grace were to take a risk acceptance strategy, she would simply acknowledge that her organization was vulnerable to SQL injection attacks and take no further action.

If Grace wanted to pursue risk avoidance, she would shut down any systems that might be vulnerable to SQL injection attacks.

If Grace wanted to pursue risk transference, she would transfer the potential financial burden of the risk to a third party, such as by purchasing an insurance policy.

Correct Answer: C. Risk mitigation

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 5 Governance Concepts: Objective 1.3

Understand Governance Concepts

Security governance processes are an integral part of safeguarding an organization's data and infrastructure. Policies, standards, procedures, frameworks, guidelines, regulations, and laws all play an important role in directing the actions of cybersecurity professionals and other team members.

In this chapter, you'll learn about CC objective 1.3. The following subobjectives are covered in this chapter:

Regulations and laws

Frameworks and guidelines

Policies, standards (e.g. International Organization for Standardization [ISO], Center for Internet Security), procedures

Laws and Regulations

Whenever we work with sensitive information, we encounter laws and regulations that govern how we collect, store, process, transmit, and dispose of that information. One of the first things that we need to figure out when working with sensitive data is what specific laws and regulations apply to us. While that might sound straightforward at first glance, the question of which jurisdictions have authority to regulate data is actually quite complicated, and compliance risks can impact an organization's risk posture.

Consider a simple example. Imagine a company with all its operations in California. It's clear in this case that California state and local laws and regulations apply to the company, as does federal law enacted at the national level. But what if the company has a customer located in New York? Does New York law now apply as well? And if the company uses a cloud provider based in Texas, does Texas law govern the data? If that cloud provider outsources to a data center provider in Florida, then what?

The issue becomes even more complicated when expanded internationally. The European Union says that its *General Data Protection Regulation*

(*GDPR*) applies to the personal information of all persons located in the EU at the time their data is being processed, regardless of their citizenship. Of course, GDPR isn't the only regulation that you'll need to follow. For example, health care providers in the United States must comply with the Health Insurance Portability and Accountability Act (HIPAA), while financial institutions must comply with the Gramm-Leach-Bliley Act (GLBA). Security professionals should be aware of the different international, national, state/territorial, and local laws that apply to their operations.

Some regulations come from sources other than the law. For example, the *Payment Card Industry Data Security Standard (PCI DSS)* is a self-regulatory program that applies to credit and debit card transactions worldwide. Compliance is enforced by the banks that provide access to the payment card system.

There's no easy answer to these jurisdictional questions. You'll need to sort through these sometimes-conflicting laws and regulations with the help of your attorneys and develop a compliance strategy appropriate to your operating environment.

Frameworks and Guidelines

Cybersecurity professionals rely on *frameworks* and *guidelines* to shape their security programs.

Frameworks provide structured approaches for managing cybersecurity risk. They typically define categories of security activities and outcomes but allow organizations flexibility in how they achieve those outcomes.


Guidelines are more specific recommendations that advise organizations on best practices for implementing security controls. While frameworks describe *what* an organization should address, guidelines often describe *how* to approach those activities. Several widely recognized frameworks and guidelines are used across the industry.

NIST Cybersecurity Framework (CSF)

The National Institute of Standards and Technology (NIST) Cybersecurity Framework (CSF) is one of the most widely adopted cybersecurity

frameworks. Originally developed for critical infrastructure organizations in the United States, it has since been adopted by organizations of all types and sizes worldwide.

The current version, CSF 2.0, organizes cybersecurity activities into six core functions: Govern, Identify, Protect, Detect, Respond, and Recover, as shown in [Figure 5.1](#). These functions provide a high-level view of the cybersecurity lifecycle and help organizations manage their cybersecurity risk effectively.



[Figure 5.1](#) NIST Cybersecurity Framework (CSF) core functions.

Source: NIST/Public domain.

NIST Risk Management Framework (RMF)

The NIST Risk Management Framework (RMF) provides organizations with a disciplined, structured process for integrating security and risk management activities into the system life cycle. While the CSF focuses on organizing an overall cybersecurity program, the RMF provides a step-by-step process for managing the risk associated with information systems and organizations.

As shown in [Figure 5.2](#), the RMF consists of seven steps. The process begins with Prepare, which establishes the context and priorities for managing

security and privacy risks across the organization. The next six steps then follow a cyclical process.




Figure 5.2 NIST risk management framework (RMF).

Source: NIST/Public domain.

The RMF steps are as follows:

- . **Prepare:** Establish the foundation and support the RMF process.
- . **Categorize:** Classify the information system and the data it processes based on their impact levels.
- . **Select:** Choose an appropriate set of security controls based on that categorization.
- . **Implement:** Put those selected controls into action within the information system.

- **Assess:** Evaluate whether the controls are in place, operating as intended, and producing the desired outcomes.
- **Authorize:** A senior official makes a risk-based decision to allow the system to operate.
- **Monitor:** Ongoing observation of the controls and the system's risk posture to ensure continued effectiveness.

This cyclical process ensures that risk management remains a continuous activity rather than a one-time event.

Security Policies, Standards, and Procedures

Security professionals do a lot of writing! We need clearly written guidance to communicate security obligations, expectations, and responsibilities to business leaders, end users, and each other. In some cases, we're setting mandatory rules that everyone in the organization must follow; in others, we're simply giving advice. Each of these roles requires communicating a bit differently.

That's where the security policy framework comes into play. Most security professionals recognize a policy framework comprising three types of documents: policies, standards, and procedures.

Security Policies

Security policies are the bedrock documents that underpin an organization's information security program. They are often developed over a long period of time and carefully written to describe an organization's security expectations.

Compliance with policies is mandatory, and policies are often approved at the very highest levels of an organization. Because of the rigor required to develop security policies, authors should strive to write them in a way that will stand the test of time.

For example, statements like "All sensitive information must be encrypted with AES-256 encryption" or "Store all employee records in Room 225" are not good policy statements. What happens if the organization switches encryption technologies or moves its records room?

Instead, a policy might state, “Sensitive information must be encrypted both at rest and in transit using technology approved by the IT department,” and “Employee records must be stored in a location approved by Human Resources.” Those statements are much more likely to stand the test of time.

Security Standards

Security standards prescribe the specific details of security controls that the organization must follow. Standards derive their authority from policy. In fact, an organization’s security policy is likely to include specific statements granting the IT department authority to create and enforce standards.

Standards are the place to include items such as the company’s approved encryption protocols, record storage locations, configuration parameters, and other technical and operational details.

Even though standards might not go through as rigorous a process as policies, compliance with them is still mandatory.

In addition to internally developed standards, organizations often adopt external standards published by recognized industry bodies. These external standards provide well-established benchmarks for security practices and can help organizations demonstrate compliance with regulatory requirements.

The International Organization for Standardization (ISO) publishes the ISO 27001 and ISO 27002 standards, which are among the most widely adopted cybersecurity standards worldwide. ISO 27001 describes the requirements for establishing, implementing, maintaining, and continually improving an information security management system (ISMS). ISO 27002 provides a detailed set of security controls that organizations can use to implement the requirements of ISO 27001.

The Center for Internet Security (CIS) publishes the CIS Controls, a prioritized set of cybersecurity best practices, and the CIS Benchmarks, which provide detailed configuration guidance for securing specific technologies. The CIS Controls are organized into three implementation groups, allowing organizations of different sizes and risk profiles to adopt a level of security appropriate to their needs.

Security Procedures

Procedures are step-by-step instructions that employees must follow when performing a specific security task. For example, the organization might have a procedure for activating the incident response team that involves sending an urgent SMS alert to team members, starting a video conference, and notifying senior management.

Compliance with procedures is mandatory.

Exam Tip

When taking the CC exam, be sure that you keep straight the differences between policies, standards, procedures, and guidelines. Specifically, remember that compliance with policies, standards, and procedures is always mandatory. Complying with guidelines is optional.

Exam Essentials

Security professionals must be aware of all the laws and regulations that apply to their organizations, including international, national, state, and local. Important examples include the European Union's General Data Protection Regulation (GDPR), the Health Insurance Portability and Accountability Act (HIPAA), the Gramm-Leach-Bliley Act (GLBA), and the Payment Card Industry Data Security Standard (PCI DSS).

Cybersecurity frameworks provide structured approaches to managing security and privacy programs. Key frameworks include the NIST Cybersecurity Framework (CSF) and the NIST Risk Management Framework (RMF).

Guidelines are more specific recommendations that advise organizations on best practices for implementing security controls. While frameworks describe *what* an organization should address, guidelines often describe *how* to approach those activities.

The common types of security governance documents include policies, standards, and procedures. Compliance with policies, standards, and procedures is mandatory. Organizations may also adopt external standards such as ISO 27001/27002 and CIS Controls/Benchmarks.

Practice Question 1

Your organization is planning to accept credit cards for the first time, and you are concerned about the regulations that may affect card processing. You already handle a large amount of personally identifiable information (PII). Which new regulation is most likely to affect your organization?

- A. GDPR
- B. PCI DSS
- C. CCPA
- D. HIPAA

Practice Question 2

You are writing a document that explains the step-by-step process that your organization's help desk should follow when helping a user reset a forgotten password. Which type of document are you creating?

- A. Policy
 - B. Standard
 - C. Procedure
 - D. Framework
-

Practice Question 1 Explanation

The Payment Card Industry Data Security Standard (PCI DSS), which regulates the storage, processing, and transmission of credit and debit card information, would most likely apply in this scenario.

The European Union's General Data Protection Regulation (GDPR) and the California Consumer Privacy Act (CCPA) may also apply to this information, depending on where the organization operates. However, this would not be a new regulation because the organization already processes other PII, and both GDPR and CCPA would apply if the company operates in the European Union and/or California.

The Health Insurance Portability and Accountability Act (HIPAA) governs protected health information (PHI). There is no indication in this question that any health records are involved.

Correct Answer: B. PCI DSS

Practice Question 2 Explanation

The key to answering this question correctly is noticing that the document is a step-by-step process. This type of process should be documented in a procedure.

Security policies are high-level documents that do not contain technical details and certainly would not include the specific steps followed by a help desk. Standards normally describe technical requirements for information and systems. Frameworks, such as the NIST Cybersecurity Framework (CSF), provide a structured approach to managing cybersecurity programs but are not step-by-step operational documents.

Correct Answer: C. Procedure

CCSM *Certified in Cybersecurity Study Guide*, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 6 Cybersecurity Controls: Objective

1.4 Understand Cybersecurity Controls

Security professionals spend most of their time designing, implementing, and managing security controls to mitigate the risks they identify during risk assessments. It's important that security professionals have a strong understanding of the different types of controls and how they work.

In this chapter, you'll learn about CC objective 1.4. The following subobjectives are covered in this chapter:

Technical controls

Administrative controls

Physical controls

What Are Cybersecurity Controls?

Cybersecurity controls are procedures and mechanisms an organization puts in place to address security risks. This might include reducing the likelihood of a risk materializing, minimizing the impact if it does, or detecting security issues that do occur.

Before moving into cybersecurity, consider how you secure your home. You probably use a variety of different security controls. Some of them might include the following:

Locks on your doors and windows designed to keep out intruders

A burglar alarm designed to detect intrusions

Security cameras to record activity inside and outside your home

Automatic light switches to deter a burglar by simulating human activity

Some of these controls are designed to achieve the same purpose or, in the language of security professionals, the same *control objective*. For example, both a burglar alarm and security cameras are designed to detect intruders. We sometimes use more than one control to achieve the same objective to ensure

we remain secure even if one control fails. If a burglar manages to open a window without tripping the burglar alarm, they may still be caught on your security cameras. This is known as the *defense in depth principle*—applying multiple overlapping controls to achieve the same objective.

Categorizing Security Controls

Security professionals use a variety of categories to group similar security controls. I'll talk about two different ways. First, I'll discuss grouping controls by their purpose: whether they are designed to prevent, detect, or correct security issues. Then I'll discuss them by their mechanisms of action: the ways they work. This groups them into the categories of technical, administrative, and physical controls.

Categorizing Controls by Purpose

When security controls are categorized by purpose, they are grouped by intent—that is, whether they are designed to prevent, detect, or correct security incidents.

Preventive controls are designed to stop a security issue from occurring in the first place. A firewall that blocks unwanted network traffic is an example of a preventive control.

Detective controls identify potential security breaches that require further investigation. This can be either during or after a breach. An intrusion detection system that searches for signs of network breaches is an example of a detective control.

Corrective controls remediate security issues that have already occurred. If ransomware infects a system and wipes out critical information, restoring that information from backup is an example of a corrective control.

Let's dig into that ransomware example a little more and talk about the ways that different controls can help protect you against ransomware. As you may know, ransomware infects systems and encrypts files, preventing users from accessing them until they pay a ransom. Consider the following three different security controls that can be put in place to protect against ransomware:

You can perform system hardening. That locks down your system's security to prevent attacks in the first place. System hardening helps prevent ransomware infections, so it is a preventive control.

You can run regular antivirus scans to detect ransomware on your system. Those scans will alert you to the presence of an infection, so they are a detective control.

If you do get a ransomware infection, those preventive and detective controls won't help you recover your data. That's where you need your backups: a corrective control.

Categorizing Controls by Mechanism of Action

The second way to categorize security controls is by their mechanism of action. This approach groups controls into technical, administrative, and physical categories.

Technical controls are exactly what the name implies, the use of technology to achieve security objectives. Think about all of the components of an IT infrastructure that perform security functions. Firewalls, intrusion prevention systems, encryption, data loss prevention, and antivirus software are all examples of technical security controls.

Exam Tip

If you see someone use the term *logical controls*, that's the same thing as *technical controls*. *Logical* and *technical* are two different words that describe the same category of control.

Administrative controls include the processes you put in place to manage technology securely. These include many of the tasks that security professionals carry out each day, such as user access reviews, policies, risk management, compliance activities, log monitoring, background checks, and security awareness training.

Physical controls are those that impact the physical world. Locks are used to keep people out of buildings, cameras to detect unauthorized intrusions, and

security guards to monitor activity in our facilities. All of these are examples of physical controls.

As you design security for your organization, you'll want to use a variety of control types and categories to help achieve your security objectives.

Exam Essentials

Security controls can be classified by purpose and mechanism of action.

The purpose types are preventive, detective, and corrective. Preventive controls are designed to stop a security issue from occurring in the first place. Detective controls identify potential security breaches that require further investigation. Corrective controls remediate security issues that have already occurred. If ransomware infects a system and wipes out critical information, restoring that information from backup is an example of a corrective control.

The mechanism of action categories are technical, administrative, and physical. Technical controls use technology to achieve security objectives. Administrative controls are the policies and processes you put in place to manage technology securely. Physical controls are those that impact the physical world.

Practice Question 1

Tonya is concerned about the risk that an attacker will attempt to gain access to her organization's database server. She is searching for a control that would block the attacker's attempts to gain access. Which type of security control is she seeking to implement?

- A. Technical
- B. Detective
- C. Corrective
- D. Preventive

Practice Question 2

Tonya evaluated all available options for protecting her database and decided to implement strong encryption to protect its contents. Which mechanism of action is she using?

- A. Technical
- B. Administrative
- C. Preventive
- D. Physical

Practice Question 1 Explanation

Tonya is attempting to stop the attacker from gaining access to a system in the first place, so this is an example of a preventive control. She hopes to prevent a security incident.

Corrective controls are used to remediate security issues that have already occurred. If an attacker compromised a system and Tonya rebuilt it to restore it, that would be a corrective control.

Detective controls identify potential security breaches. If Tonya installed an intrusion detection system that alerted her to compromised systems, that would be an example of a detective control.

It is possible that Tonya may deploy a technical control, but all we know from this scenario is the *purpose* of the desired control: to prevent compromises. No mechanism of action is specified for that preventive control.

Correct Answer: D. Preventive

Practice Question 2 Explanation

Technical controls use technology to achieve security objectives. In this scenario, Tonya is using encryption to protect the contents of her database, thereby implementing a technical control.

Administrative controls are policies and processes put in place to manage technology securely. There is no discussion of processes in this scenario, so Tonya is not implementing an administrative control.

Physical controls are those that impact the physical world. Encryption is a digital technology and does not impact the physical world, so it is not a physical control.

This use of encryption is a preventive control, but preventive is a purpose type, not a mechanism of action category. The question specifically asked for the mechanism of action.

Correct Answer: A. Technical

CCSM Certified in Cybersecurity Study Guide, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 7 Professional and Ethical Conduct:

Objective 1.5 Maintain Professional and Ethical Conduct

Ethics guide how we behave and handle our personal and professional responsibilities. While each of us has our own internal sense of ethics, we are also subject to ethical codes provided by others. Many employers have written ethical standards that apply to all employees. ISC2 has both a Code of Ethics and a Code of Professional Conduct that apply to the behavior of all ISC2 members and Associates.

In this chapter, you'll learn about CC objective 1.5. The following subobjectives are covered in this chapter:

Professional code of conduct

Due care and due diligence

ISC2 Code of Ethics

Corporate Codes of Ethics

Corporate codes of ethics require that information security professionals act honorably and responsibly. Because information security professionals are often entrusted with highly sensitive data, it is critical that they act ethically and in a trustworthy manner.

Many organizations have internal codes of ethics that their employees must follow. These codes outline the principles and guidelines that employees are expected to follow to act honestly and ethically and avoid personal conflicts of interest.

For example, [Figure 7.1](#) shows a portion of the Code of Business Conduct for AT&T. It includes rules and guidelines for how employees should conduct themselves honestly and ethically and avoid personal conflicts of interest.



We act with integrity every day.

At AT&T, we act with integrity every day. We:

- **Follow the Code of Business Conduct (the "Code") and all policies.** We comply with the letter and spirit of AT&T's Code and policies. Our responsibility to act ethically extends beyond our day-to-day tasks. This can include actions that occur outside of work, especially those that might affect our job performance or impact the Company's reputation.
- **Speak up.** We **report** any suspected violations of the Code, our policies, or any legal or regulatory requirements.
- **Are accountable for our actions.** Any failure to comply with this Code, our policies, or any legal or regulatory requirements may lead to disciplinary action, up to and including termination of employment. No one has the authority to instruct any employee to break the law, violate this Code, or infringe upon our policies.
- **Use good judgment and ask questions.** The Code serves as a robust foundation for ethical business conduct. However, it is not a replacement for sound judgment. When the right course of action is unclear, we have [AT&T's Ethical Decision-Making Model](#) to guide us.

Figure 7.1 AT&T's code of business conduct.

Source: AT&T Intellectual Property / Code of Business Conduct /

<https://about.att.com/content/dam/pages/ethics/COBC.pdf> / last accessed Apr. 30, 2026

AT&T's code is quite typical of those used by large organizations. If you'd like to review this multipage document in detail, you can find it at

<https://about.att.com/content/dam/pages/ethics/COBC.pdf>.

By adhering to corporate codes of ethics, information security professionals can build trust with their colleagues and clients and protect themselves from legal and ethical repercussions. These codes promote a culture of transparency, accountability, and responsibility within organizations, ultimately benefiting everyone involved.

Due Care and Due Diligence

Information security professionals have a legal and ethical obligation to act responsibly when protecting their organization's information and systems. Two important concepts describe this obligation: due care and due diligence.

Although they are closely related, they refer to different aspects of responsible behavior.

Due Care

Due care means doing what a reasonable person would do in a given situation to protect the organization. It is about taking the right actions. An organization exercises due care when it implements security policies, deploys appropriate controls, trains its employees on security awareness, and responds promptly to known vulnerabilities and incidents. For example, an organization that installs firewalls, requires strong passwords, and encrypts sensitive data is demonstrating due care by taking reasonable steps to protect its assets.

The legal concept of due care is sometimes described as the “prudent person” standard. If a security breach occurs and the organization can demonstrate that it took reasonable precautions, it may reduce its legal liability. On the other hand, an organization that fails to implement basic security measures has failed to exercise due care and may face legal consequences.

Due Diligence

Due diligence is the ongoing effort to investigate, monitor, and maintain an organization’s security protections. While due care is about taking action, due diligence is about doing the homework needed to ensure those actions are effective. An organization demonstrates due diligence by conducting risk assessments, performing vulnerability scans, auditing its security controls, reviewing access logs, and staying informed about emerging threats.

Think of the relationship this way: due care is implementing a security policy that requires regular system patching, while due diligence is verifying that those patches are actually applied on schedule. Both concepts work together. An organization that writes security policies but never checks whether they are being followed has exercised due care but not due diligence. An organization that monitors its systems but has no formal security policies in place may be exercising due diligence without due care.

Exam Tip

Remember the distinction between due care and due diligence. Due care is about taking reasonable actions to protect the organization (“doing the right thing”), while due diligence is about the ongoing effort to verify and maintain those protections (“doing the homework”). Exam questions may present scenarios and ask you to identify which concept applies.

ISC2 Code of Ethics

ISC2 has a *code of ethics* that applies to all ISC2 members and Associates. The Code sets forth a clear purpose in its preamble:

The safety and welfare of society and the common good, duty to our principals, and duty to each other, require that we adhere, and be seen to adhere, to the highest ethical standards of behavior.

Therefore, strict adherence to this Code is a condition of certification.

The Code then enumerates four canons: four simple statements that outline what is expected of individuals who subscribe to the Code:

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

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

Provide diligent and competent service to principals.

Advance and protect the profession.

You’ll want to be very familiar with these canons when you take the exam. You don’t necessarily need to be able to recite them word-for-word, but you do need to know the general idea behind each one.

Canon I

The first canon of the Code is that you must protect society, the common good, and its infrastructure. The actions you take should give the public confidence in our profession. This canon basically means that the actions you take, or fail to take, must support the betterment of society. As a certified security professional, you have an obligation to protect the common good. For example, a security professional who uses their skills to engage in unethical hacking activities would violate this canon of the Code.

Canon II

The second canon is that your actions must be ethical. You must act with honor, justice, and responsibility and work within the bounds of the law. The bottom line here is that you may not break the law, lie, or commit any other dishonorable, unjust, or irresponsible action. For example, a security professional who makes an error that leads to a compromise at their organization, then covers it up and lies about the mistake, violates this canon of the ISC2 Code of Ethics.

Canon III

The third canon is that the professional services you provide to principals must be both diligent and competent. As a security professional, you must carry out your duties responsibly. The Code uses the word *principal* here because it applies to your employer if you are a normal employee or to your clients if you are a consultant. Basically, whoever you are working for has the right to expect your diligent and competent service. A security professional who fails to fulfill their assigned and agreed-upon duties violates this canon.

Canon IV

Finally, the fourth canon is that your actions as a security professional should advance and protect the information security profession. The actions you take should help, rather than detract from, the profession at large. The most common way that this canon is violated is when certified individuals provide unauthorized assistance on exams, violate the ISC2 nondisclosure agreement, or provide false information on an applicant's endorsement application.

Exam Tip

When you take the CC exam, you should not only remember the main idea behind each of the four canons, but you should also be able to analyze an ethical issue and identify the canon(s) that might have been violated.

Ethical Complaint Procedures

You must follow a formal process if you suspect that another ISC2 member has violated the ISC2 Code of Ethics. The code includes a provision requiring you to report violations. In fact, failure to report a violation that you know of is itself a violation of the code of ethics. So, if you witness a violation of the code of ethics and do not report it, you yourself are in violation of the code and subject to sanctions.

If you need to submit a violation report, you must provide a written, notarized affidavit using the form available on the ISC2 website. It must include the name of the accused person, the nature of the violation, the specific canon or canons breached, the reason you have standing, and any corroborating evidence.

Having *standing* to file a complaint means that the alleged behavior must harm you or your profession in some way. Standing varies based on the canon involved.

The first two canons are about protecting society at large and acting responsibly. Anyone may be harmed by those violations, so any member of the public has standing to file a complaint about the first two canons.

The third canon concerns service to principals (employers or clients), so only employers or clients of the individual have standing to file a complaint under it.

The last canon concerns protecting the profession, so other professionals have standing to file a complaint about violations of the last canon. This doesn't mean that you have to be ISC2 certified or even a security professional. Anyone who is certified or licensed in any field and subscribes to a code of

ethics may file a complaint with this canon. For example, a certified public accountant or a licensed health care professional would have standing to file a complaint with this canon.

Once a complaint is filed, the ISC2 Ethics Committee takes over. They will allow the accused individual to respond, gather any additional evidence that they wish, and reach a determination. If they find an individual has violated the Code of Ethics, they may revoke that individual's certification.

Code of Professional Conduct

In addition to the Code of Ethics, ISC2 has a Code of Professional Conduct, a global framework that establishes clear expectations for the responsibilities and obligations of cybersecurity leaders and practitioners. Built on the foundation of the ISC2 Code of Ethics, the Code of Professional Conduct provides more detailed, actionable guidance on the ethical challenges cybersecurity professionals face in their daily work. It is meant to supplement, not to replace, the ISC2 Code of Ethics.

The Code of Professional Conduct is organized around two guiding principles: ethics and professional conduct.

The ethics principle addresses topics including integrity, confidentiality, respect for laws and regulations, and public safety and societal impact. This principle includes the following commitments:

- Commitment to honesty and transparency

- Commitment to the profession

- Commitment to responsible use of privileges

- Commitment to data security and privacy

- Commitment to compliance and your organization's code of conduct

- Commitment to security and risk reduction

- Commitment to responsibly securing emerging technologies

The professional conduct principle covers responsibility and accountability, collaboration and teamwork, competence and continuous improvement, and reporting issues and concerns. It includes the following commitments:

Commitment to balancing security and organizational objectives

Commitment to employers, employees, and stakeholders

Commitment to respect for others

Commitment to ethical leadership & workplace culture

Commitment to competence and continuous improvement

Commitment to ethical reporting

Unlike the ISC2 Code of Ethics, which applies specifically to ISC2 members and Associates, the Code of Professional Conduct is intended as a resource for the entire cybersecurity community, regardless of certification or organizational affiliation. The Code helps practitioners navigate complex situations involving data privacy, the balance between security and organizational objectives, incident response decisions, AI governance, and emerging threats.

Exam Tip

It's really important that you understand who can file a complaint under each canon. You should be able to read a scenario, recognize which canons might have been violated, and describe who has standing to file a complaint under that canon.

Exam Essentials

The first canon of the ISC2 Code of Ethics says that you must protect society, the common good, necessary public trust and confidence, and the infrastructure. Anyone may file a complaint against a member under the first canon.

The second canon of the ISC2 Code of Ethics says that you must act honorably, honestly, justly, responsibly, and legally. Anyone may file a complaint against a member under the second canon.

The third canon of the ISC2 Code of Ethics says that you must provide diligent and competent service to principals. Only one of those principals (an individual's employer or consulting client) may file a complaint against a member under the third canon.

The fourth canon of the ISC2 Code of Ethics says that you must advance and protect the profession. Any certified professional who subscribes to an ethical code may file a complaint against a member under the fourth canon.

Due care is the responsibility to take reasonable actions to protect the organization, such as implementing security policies and deploying controls. Due diligence is the ongoing effort to investigate, monitor, and verify that those protections are working effectively. Both concepts work together to demonstrate responsible security management.

The ISC2 Code of Professional Conduct provides a global framework for ethical and professional behavior in cybersecurity. It is organized around two guiding principles: ethics and professional conduct. It covers topics including integrity, confidentiality, respect for laws, public safety, responsibility and accountability, collaboration, competence, and reporting issues. Unlike the Code of Ethics, it is intended for all cybersecurity professionals, not just ISC2 members.

Practice Question 1

You are the supervisor of a team of cybersecurity professionals, all of whom hold current ISC2 certifications. You believe that one of those employees has stolen funds from your organization, and an internal investigation confirmed the likelihood of that action. Which two canons of the Code of Ethics were most likely violated?

- A. Canons I and II
 - B. Canons II and III
 - C. Canons III and IV
 - D. Canons I and IV
-

Practice Question 2

You have witnessed behavior by a certified cybersecurity professional and ISC2 member that you believe directly harms the cybersecurity profession. Who has standing to file a complaint about this behavior?

- A. Any member of the public
 - B. Any certified professional
 - C. The individual's employer or consulting client
 - D. Nobody
-

Practice Question 1 Explanation

Canon I is that you must protect society, the common good, necessary public trust and confidence, and the infrastructure. Although stealing from an employer is certainly unethical, it likely doesn't rise to the level of harming society described in this question.

Canon II is that you must act honorably, honestly, justly, responsibly, and legally. In this case, the employee who stole certainly did not act in accordance with these principles. Based on this information alone, you can select the answer "Canons II and III" because it is the only one that includes Canon II, but let's continue the analysis to confirm.

Canon III requires you to provide diligent and competent service to principals (such as your employer). Stealing from your employer is certainly not competent service, so there is a good argument that Canon III was violated.

Canon IV requires you to advance and protect the profession. As with the Canon I analysis, the employee's behavior was certainly unethical, but it doesn't rise to the level of harming the profession.

Correct Answer: B. Canons II and III

Practice Question 2 Explanation

To answer this question, you must first identify which canon of the ISC2 Code of Ethics is involved. Canon IV of the code is to “Advance and protect the profession,” which seems to match the scenario provided here.

Next, you must know who has standing to file a complaint under that particular canon. According to the complaint procedures, any certified professional may file a complaint under Canon IV.

Any member of the public is entitled to file an ethics complaint under Canons I and II but not under Canon IV.

An individual’s employer or consulting client may file an ethics complaint under Canon III but not under Canon IV.

Correct Answer: B. Any certified professional

CCSM Certified in Cybersecurity Study Guide, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Part II Domain 2: Security Governance

[Chapter 8 Governance, Risk, and Compliance](#)

[Chapter 9 Business Continuity](#)

[Chapter 10 Disaster Recovery](#)

[Chapter 11 Security Awareness](#)

[Chapter 12 Measuring Cybersecurity Effectiveness](#)

Security Governance is the second domain of ISC2's Certified in Cybersecurity exam. It provides the information that cybersecurity professionals need to know to plan and implement governance programs, ensure the continuity of operations, build security awareness programs, and measure security effectiveness. This domain includes the following four objectives:

2.1 Plan Governance, Risk, and Compliance (GRC)

2.2 Understand Redundancy

2.3 Understand Security Awareness

2.4 Measure Cybersecurity Effectiveness

Questions from this domain make up approximately 17% of the questions on the CC exam.

CCSM Certified in Cybersecurity Study Guide, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 8 Governance, Risk, and Compliance:

Objective 2.1 Plan Governance, Risk, and Compliance (GRC)

In today's complex business environment, organizations must juggle an ever-growing set of security responsibilities, from protecting sensitive systems to meeting regulatory and audit requirements. Managing these responsibilities in an ad hoc fashion quickly becomes unsustainable, and organizations need a structured approach that ties together their security decision-making, threat management, and regulatory obligations. That's where *governance, risk, and compliance* come in.

In this chapter, you'll learn about CC objective 2.1. The following subobjectives are covered in this chapter:

Purpose

Importance

Frameworks and tools

Purpose of Governance, Risk, and Compliance

Organizations face an ever-growing array of security challenges, regulatory requirements, and strategic decisions. *Governance, risk, and compliance (GRC)* is an integrated approach that brings together three closely related disciplines to help organizations achieve their objectives while managing uncertainty and meeting their obligations. Rather than treating governance, risk management, and compliance as separate silos, GRC unifies them into a coordinated strategy that aligns security activities with business goals.

Understanding Governance

Governance is the system of policies, processes, and structures that an organization's leadership uses to direct and control the organization. In the context of cybersecurity, governance ensures that security activities support

the organization's mission and objectives. It defines who has authority to make security decisions, how those decisions are made, and how accountability is maintained.

As discussed in [Chapter 5, “Governance Concepts,”](#) governance relies on a hierarchy of documents, including policies, standards, and procedures that communicate security expectations throughout the organization. While [Chapter 5](#) focused on the specific types of governance documents and their characteristics, this chapter examines how governance fits into the broader GRC strategy and how organizations plan and implement effective governance programs.

Understanding Risk Management

Risk management is the ongoing process of identifying, assessing, prioritizing, and responding to risks that could affect an organization's ability to achieve its objectives. In cybersecurity, risks arise from threats that could exploit vulnerabilities in systems, processes, dependencies, physical environments, or people. Effective risk management helps organizations make informed decisions about where to invest their limited security resources.

As you learned in [Chapter 4](#), risk management process typically follows a lifecycle that includes identifying, assessing, treating, and monitoring risks. Organizations may choose to mitigate risks by implementing controls, transferring risks through insurance or contracts, avoiding risks by eliminating the activity that creates them, or accepting risks when the cost of treatment exceeds the potential risk.

Understanding Compliance

Compliance refers to an organization's adherence to laws, regulations, standards, and internal policies that govern its operations. As discussed in [Chapter 5](#), organizations must navigate a complex landscape of regulations and industry standards, including the European Union's General Data Protection Regulation (GDPR), the Payment Card Industry Data Security Standard (PCI DSS), and various national and local laws. Compliance programs ensure that the organization systematically meets these obligations and can demonstrate adherence to auditors and regulators.

Integrating Governance, Risk, and Compliance

The true power of GRC lies in integrating the three disciplines of governance, risk, and compliance. When governance, risk management, and compliance operate independently, organizations often experience duplication of effort, conflicting priorities, and gaps in coverage. An integrated GRC approach connects these activities so that governance decisions are informed by risk assessments, compliance requirements shape risk priorities, and risk management outcomes feed back into governance policies.

For example, when a new regulation takes effect, an integrated GRC program ensures that the compliance team identifies the new requirements, the risk team assesses the impact of noncompliance, and the governance team updates policies and standards accordingly. This coordinated response is far more efficient and effective than having each team work independently.

Importance of GRC

GRC programs are essential for organizations of all sizes. As cybersecurity threats continue to evolve and regulatory requirements become more complex, organizations need a structured approach to managing these challenges. Without an effective GRC program, organizations risk security breaches, regulatory penalties, operational disruptions, and reputational damage.

Business Alignment

One of the most important benefits of GRC is that it aligns security activities with business objectives. Security programs that operate in isolation from the business often either fail to protect what matters most or impose unnecessary restrictions that hinder productivity. GRC provides a framework for making security decisions in line with business priorities, ensuring that security investments support the organization's strategic goals.

For instance, a GRC program helps an organization determine how much risk it is willing to accept in pursuit of business opportunities. This concept, known as risk appetite, is a governance decision that guides risk management activities throughout the organization. By establishing a clear risk appetite,

leadership provides security teams with the guidance they need to make appropriate risk-based decisions.

Regulatory Compliance

Organizations that fail to comply with applicable regulations face significant consequences, including financial penalties, legal liability, and loss of the ability to operate in certain markets or industries. GRC programs help organizations systematically identify their compliance obligations, implement the necessary controls, and maintain evidence of compliance.

An effective GRC program also helps organizations prepare for audits and regulatory examinations. By maintaining organized records of policies, risk assessments, and compliance activities, organizations can demonstrate their due diligence to auditors and regulators, reducing the likelihood of adverse findings.

Operational Efficiency

When governance, risk management, and compliance activities are siloed, organizations often end up performing redundant assessments, maintaining duplicate documentation, and creating conflicting requirements. An integrated GRC approach eliminates this duplication by establishing common processes, shared risk registers, and unified reporting. This not only reduces costs but also frees security professionals to focus on higher-value activities.

Exam Tip

Remember that GRC is about integration. The exam may test your understanding of why organizations benefit from combining governance, risk management, and compliance into a unified program rather than treating them as separate activities. Key benefits include business alignment, regulatory compliance, and operational efficiency.

Stakeholder Confidence

A well-implemented GRC program demonstrates to stakeholders, including customers, partners, investors, and regulators, that the organization takes security seriously and manages its risks responsibly. This confidence can be a competitive advantage, particularly in industries where trust is a critical factor in business relationships.

Continuous Improvement

GRC programs provide the feedback mechanisms needed for continuous improvement. By monitoring compliance status, tracking risk trends, and reviewing the effectiveness of governance structures, organizations can identify areas for improvement and adapt their programs to address emerging threats and changing business conditions.

Frameworks and Tools

Organizations rely on established frameworks and specialized tools to implement and manage their GRC programs. Frameworks provide the structure and methodology for GRC activities, while tools automate and streamline day-to-day processes for governance, risk, and compliance management.

GRC Frameworks

In [Chapter 5](#), I discussed several important cybersecurity frameworks, including the NIST Cybersecurity Framework (CSF) and the NIST Risk Management Framework (RMF). These frameworks play an important role in GRC programs, but there are additional frameworks specifically designed to address the broader scope of governance, risk, and compliance. These include:

COBIT is a framework developed by ISACA for the governance and management of enterprise information and technology. COBIT provides a comprehensive set of controls and processes that help organizations ensure their IT activities align with business objectives, manage IT-related risks, and

comply with relevant regulations. COBIT is particularly useful for organizations seeking to establish or improve their IT governance practices.

ISO 31000 is an international standard that provides principles and guidelines for risk management. Unlike frameworks that focus specifically on cybersecurity or IT risks, ISO 31000 applies to all types of organizational risk. It establishes a common risk management vocabulary and process that organizations can adapt to their specific needs. The standard emphasizes that risk management should be integrated into all organizational activities, including strategic planning, decision-making, and operational processes.

ISO 27001 is an international standard for information security management systems (ISMS) that covers information security, cybersecurity, and privacy. As noted in [Chapter 5](#), ISO 27001 describes the requirements for establishing, implementing, maintaining, and continually improving an ISMS. In the context of GRC, ISO 27001 provides a systematic approach to managing information security risks and ensuring compliance with security requirements.

GRC Tools and Platforms

As GRC programs grow in complexity, organizations increasingly rely on technology solutions to manage their governance, risk, and compliance activities. GRC platforms are integrated software solutions that provide a centralized environment for managing all aspects of GRC. They often include features such as:

Risk Registers. A risk register is a fundamental GRC tool that documents identified risks, including their likelihood, impact, risk response strategies, and current status. Risk registers provide a centralized view of an organization's risk landscape and help risk managers track risks over time. Many organizations maintain risk registers in spreadsheets, but dedicated GRC platforms offer more sophisticated risk tracking and reporting capabilities.

Compliance Tracking. These tools help organizations monitor their compliance status across multiple regulations, standards, and internal policies. Compliance tracking systems typically map controls to specific requirements, track each control's implementation status, and generate reports that demonstrate compliance to auditors and regulators.

Policy Management. Policy management tools automate the lifecycle of governance documents, including creation, review, approval, distribution, and retirement. These tools ensure that policies remain current, that employees acknowledge their responsibility to follow them, and that version history is maintained for audit purposes.

Audit Management. Audit management tools help organizations plan, execute, and track internal and external audits. They facilitate the collection of audit evidence, management of findings, and tracking of remediation activities. By automating the audit process, these tools help organizations maintain a continuous state of audit readiness.

Note

When selecting GRC frameworks and tools, organizations should consider their specific industry, regulatory environment, and organizational size and maturity level. There is no one-size-fits-all solution, and many organizations use a combination of frameworks and tools tailored to their unique needs.

Exam Essentials

Understand the purpose of governance, risk, and compliance (GRC). GRC is an integrated approach that combines governance, risk management, and compliance into a coordinated strategy. Governance provides the policies and structures for decision-making, risk management identifies and responds to potential threats, and compliance ensures adherence to laws, regulations, and internal policies.

Know why GRC is important to organizations. GRC programs align security activities with business objectives, help organizations meet regulatory requirements, improve operational efficiency by eliminating silos, build stakeholder confidence, and enable continuous improvement of security programs.

Be familiar with GRC frameworks and tools. Key GRC frameworks include COBIT for IT governance, ISO 31000 for risk management, and ISO 27001 for information security management systems. The NIST Cybersecurity Framework (CSF) and Risk Management Framework (RMF) also support GRC activities. Core GRC tools include risk registers, compliance tracking systems, policy management tools, and audit management systems.

Practice Question 1

Your organization currently manages governance, risk management, and compliance as three separate programs, each with its own team, documentation, and reporting structure. Leadership is considering integrating these into a unified GRC program. Which of the following is the most significant benefit of this integration?

- A. It eliminates the need for compliance audits.
 - B. It reduces duplication of effort and aligns security activities with business objectives.
 - C. It transfers all organizational risk to a third party.
 - D. It replaces the need for security policies and standards.
-

Practice Question 2

A security analyst is evaluating frameworks to help establish an IT governance program for her organization. Which of the following frameworks would be most appropriate for this purpose?

- A. PCI DSS
 - B. ISO 31000
 - C. COBIT
 - D. GDPR
-

Practice Question 1 Explanation

The most significant benefit of integrating governance, risk management, and compliance into a unified GRC program is that it reduces duplication of effort and aligns security activities with business objectives. When these disciplines operate in silos, organizations often perform redundant assessments and maintain duplicate documentation.

GRC integration does not eliminate the need for compliance audits (A), as audits remain an important verification mechanism. It does not transfer all risk to a third party (C); that would be risk transfer, which is a specific risk response strategy, not a benefit of GRC integration. GRC programs rely on security policies and standards (D) and do not replace them.

Correct Answer: B. It reduces duplication of effort and aligns security activities with business objectives.

Practice Question 2 Explanation

COBIT is specifically designed for the governance and management of enterprise IT. It provides a comprehensive framework of controls and processes that align IT activities with business objectives.

PCI DSS (A) is a compliance standard for organizations that handle payment card data, not an IT governance framework. ISO 31000 (B) is a risk management standard that applies to all types of organizational risk; it is valuable for risk management but is not specifically focused on IT governance. GDPR (D) is a European Union regulation governing data protection and privacy, not a governance framework.

Correct Answer: C. COBIT

Chapter 9 Business Continuity: Objective 2.2 Understand Redundancy

Cybersecurity professionals are responsible for maintaining the availability of systems and data as a key function. They must ensure that authorized individuals have access to the resources they need to provide a service or to get their work done. Business continuity comprises the core activities that help meet this goal.

In this chapter, you'll learn about one of the subobjectives of CC objective 2.2. The remaining material for this objective is covered in [Chapter 10, "Disaster Recovery."](#) The following subobjective is covered in this chapter:

Business Continuity (BC)

Business Continuity Planning

Business continuity planning (BCP) is a core responsibility of the information security profession. *Business continuity* efforts are activities designed to keep a business running in the face of adversity. This may come in the form of a small-scale incident, such as a single system failure, or a catastrophic environmental event, such as an earthquake or tornado. Business continuity plans may also be activated by human-caused hazards, such as a terrorist attack or a hacker intrusion.

The focus of business continuity is keeping operations running, so BCP is sometimes referred to as continuity of operations planning (COOP).

While many organizations place the responsibility for business continuity with operational engineering teams, business continuity is a core security concept because it is the primary control that supports the information security objective of availability. Remember, that's one of the "big three" goals of information security: confidentiality, integrity, and availability.

BCP Scope Definition

When an organization begins a business continuity effort, it's easy to quickly become overwhelmed by the many possible scenarios and controls that the project might consider. For this reason, the team developing a business continuity plan should take time up front to carefully define its scope. Here are three questions that can help define the scope:

Which business activities will the plan cover?

What types of systems will the plan cover?

What types of controls will the plan consider?

The answers to these questions will help you make critical prioritization decisions in later stages.

Business Impact Analysis

Continuity planners use a process known as *business impact analysis (BIA)* to inform scope decisions. The BIA is an impact assessment that begins by identifying the organization's mission-essential functions, then traces them backward to identify the critical IT systems that support those processes. Once planners have identified the affected IT systems, they identify the potential risks, including likelihood and impacts to those systems, as part of a risk assessment.

The output of a business impact analysis is a prioritized list of risks and potential impacts that might disrupt the organization's operations, as shown in [Table 9.1](#). Planners can then use this information to select controls that mitigate the organization's risks within acceptable expense limits. For example, notice that the risks in this scenario are listed in descending order of expected loss.

TABLE 9.1

Example of Prioritized Risks and Potential Impacts

Risk	Annualized Loss Expectancy
Hurricane damage to the data center	\$145,000
Fire in data center	\$18,000
Power outage	\$12,000
Theft of equipment	\$3,400

It makes sense to prioritize the highest-risk item at the top of the list (hurricane damage to the data center). However, the organization must also consider costs when deciding on control implementation. For example, if a \$50,000 flood-prevention system would reduce the risk of hurricane damage to the data center by 50%, purchasing the system is clearly a good decision because its expected payback period is less than one year.

In a cloud-centric environment, BCP becomes a collaboration between the cloud service provider (CSP) and the customer. For example, the risk of a hurricane damaging a data center may be mitigated by the service provider installing a flood-prevention system, but it may also be mitigated by the customer replicating services across data centers, availability zones, and geographic regions.

Business Continuity Controls

Business continuity professionals have a variety of tools at their disposal to help remediate potential availability issues. One of the critical ways IT professionals protect system availability is by ensuring systems are redundant and fault-tolerant. That simply means they are designed so that the failure of a single component doesn't bring down the entire system—business can continue in the face of a single predictable component failure.

Single Point of Failure Analysis

The single point of failure (SPOF) analysis process provides IT professionals with a mechanism to identify and remove single points of failure from their systems or processes.

Consider a simple example. [Figure 9.1](#) shows a simple web-based application: a web server protected by a firewall and connected to the Internet.



[Figure 9.1](#) Web-based application.

When conducting a SPOF analysis, you might first notice that the web server itself is a single point of failure. If anything goes wrong with the server, the web service will stop functioning. You can resolve this issue by replacing the single web server with a clustered web server farm, as shown in [Figure 9.2](#).



[Figure 9.2](#) Adding clustered web servers.

The cluster is designed so that if a single server fails, the remaining servers will continue to provide service without disruption. Once you've implemented the cluster, you've removed the server as a single point of failure. Next, you might turn your attention to the firewall, another single point of failure. If the firewall goes down, Internet users will not be able to reach the web server, rendering the web service unavailable. Therefore, the firewall is also a single point of failure.

You can correct this situation by replacing the firewall with a pair of high-availability firewalls, as shown in [Figure 9.3](#).



Figure 9.3 Adding high availability firewalls.

In this approach, one firewall serves as a backup device, standing by to step in immediately if the primary firewall fails. By replacing the single firewall with a high-availability pair, you've removed the firewall as a single point of failure.

There are two other single points of failure here: the internal and external network connections. As with the web server and firewall, you can address these single points of failure by introducing redundancy, as shown in [Figure 9.4](#). In this approach, each link has two separate network connections. If one fails, the service will continue to operate over the other.



Figure 9.4 Adding redundant network links.

This SPOF analysis may continue, identifying and remediating issues until either the team stops finding new issues or the cost of addressing them outweighs the potential benefit. SPOF analysis is an important part of an organization's continuity of operations planning efforts.

Other Continuity Risks

Organizations should also consider the other risks facing their IT operations. As they conduct IT contingency planning, they should consider not only single points of failure but also other situations that might jeopardize business continuity. For example, these might include:

Sudden bankruptcy of a key vendor

Inability to provide the requisite computing or storage capacity

Utility service failures

Any other risk that IT management believes may disrupt operations

One final component of business continuity planning often overlooked is personnel succession planning. Information technology depends on highly skilled team members who develop, configure, and maintain systems and processes. IT leadership should work with their human resources department to identify team members essential to continued operations and identify potential successors to those positions. That way, when someone leaves the organization, management has already thought through potential replacements and, hopefully, provided those successors with the professional development opportunities they need to step into the departing employee's shoes.

High Availability and Fault Tolerance

I've already discussed some ways that information security professionals can ensure the continued operation of systems. This section digs into this in a little more detail. There are two key technical concepts that improve system availability: high availability and fault tolerance.

High Availability

High availability (HA) uses multiple systems to protect against failures. These are techniques like those discussed in the previous section. One example is to deploy a cluster of web servers that can continue to operate even if a single server fails. Another example is to use a pair of firewalls, with one designated as the backup. The core concept of high availability is having operationally redundant systems, sometimes at different sites. The geographic dispersal of placing systems in different locations protects you against damage to a facility.

Load Balancing

Load balancing is a related but different concept from high availability. Load balancing distributes the workload across multiple systems, providing a scalable computing environment. While they use similar technologies, load balancing and high availability are different goals. Most implementations of clustering and similar technologies are designed to achieve both high availability and load balancing, but it is possible to have one without the other.

Fault Tolerance

Fault tolerance (FT) helps to protect a single system from failing in the first place by making it resilient in the face of technical failures. Three of the most common points of failure within a computer system are the device's:

Power supply

Storage media

Networking components

Fault tolerance controls can prevent the system from failing even if one of these components fails completely.

Power Supplies

Power supplies contain moving parts and, therefore, are common points of failure. If a power supply fails, the results can be catastrophic. For this reason, server manufacturers often build dual power supplies into their servers, as shown in [Figure 9.5](#).



Figure 9.5 Server with dual power supplies.

Source: SasaStock/Adobe Stock Photos

When a customer installs the server, they connect both of the power supplies to a power source. This way, if one power supply fails, the other can keep the server running uninterrupted. For added redundancy, data centers with two separate power sources can connect each power supply to a different source.

Data centers also use *uninterruptible power supplies (UPSs)*, such as the one shown in [Figure 9.6](#), to provide battery power to systems during brief disruptions. These power sources may also be served by a generator that provides long-term backup power.



Figure 9.6 Uninterruptible power supply (UPS).

Source: Sergey Ryzhov/Adobe Stock Photos

Managed *power distribution units (PDUs)* control and distribute power within a rack, ensuring the power delivered to devices is clean and controlled. Additionally, PDUs may support remote management.

Storage


The second priority of many fault tolerance efforts is to protect against the failure of a single storage device. This is achieved through the use of a technology known as *redundant arrays of independent disks (RAID)*. RAID comes in several forms, but each is designed to provide redundancy by using more disks than needed to meet business needs. Let's look at two RAID technologies: disk mirroring and disk striping with parity.

Disk Mirroring

The simplest form of RAID, RAID Level 1, is *disk mirroring*. In this approach, the server contains two disks. Each disk has identical contents, and

when the system writes any data to one disk, it automatically makes the same changes to the other disk, keeping it as a synchronized copy, or mirror, of the primary disk. If the primary disk fails, the system can automatically switch to the backup disk and continue operating.


You can see this in [Figure 9.7](#). Each of the data blocks A1 through A4 is stored on both disk 0 and disk 1, so that the two disks are identical.



[Figure 9.7](#) RAID 1 disk mirroring.

Disk Striping with Parity

The second major RAID technology is *disk striping with parity*, known as RAID 5. In this approach, the system contains three or more disks and writes data across all of them, with additional elements known as parity blocks spread across the disks, as shown in [Figure 9.8](#). If one of the disks fails, the system can regenerate its contents by using that parity information.



[Figure 9.8](#) RAID 5 disk striping with parity.

Examining [Figure 9.8](#), you can see that sections of each block are stored on two of the three disks, and the third disk contains the parity information for

that block. For example, pieces of block A (A1 and A2) are stored on disks 1 and 2, while disk 0 contains the parity information for that block (Ap). Similarly, pieces of block B (B1 and B2) are stored on disks 0 and 1, with the parity information (Bp) stored on disk 2.

Exam Tip

One important thing to remember: RAID is a fault tolerance strategy. It is *not* a backup strategy. You should still perform regular data backups to protect your organization's information in the event of a more catastrophic failure, such as the physical destruction of the entire server.

Networking Components

Networking components can also be single points of failure. Therefore, organizations should consider implementing redundancy at different points in the network. This ranges from having multiple Internet service providers enter a facility to using dual network interface cards in critical servers, similar to how we use multiple power supplies. Using two network interface cards is known as *NIC teaming*.

Within a network, add redundancy to critical network paths as well. For example, the connection between servers and their storage is crucial to the data center's operation. *Multipath* approaches create redundancy in these paths and ensure continuous access to storage.

Redundancy Through Diversity

Finally, at a higher level, think about adding diversity to your environment wherever possible to avoid redundant elements all falling victim to the same flaw at the same time.

Use diverse technologies from a diverse set of vendors to avoid a single technology or vendor failure from critically impacting your environment. You should also consider diversifying your cryptography and other security controls.

Exam Essentials

Business continuity efforts are activities designed to keep a business running in the face of adversity. The business impact analysis (BIA) is an impact assessment that begins by identifying the organization's mission-essential functions and then tracing those backward to identify the critical IT systems that support those processes.

Single point of failure analyses identify places where the failure of one component could cause an entire system or service to become unavailable.

High availability (HA) uses multiple systems to protect against failures. Fault tolerance (FT) helps to protect a single system from failing in the first place by making it resilient in the face of technical failures.

Three of the components most likely to fail in a computer system are the power supply, storage, and network connection.

Practice Question 1

Renee would like to add fault tolerance to the storage capacity of a new server that she is building. She would like to use RAID 5 to protect this server. What is the minimum number of physical disks that she can use?

- A. 1
- B. 2
- C. 3
- D. 5

Practice Question 2

Kevin is designing a new web server environment for his organization. He believes that the service will need three servers to support normal traffic levels and decides to use three servers to meet that need. As requests come in, they will be sent to the server that has the most available capacity.

Which term best describes what Kevin is doing?

- A. Fault tolerance
- B. High availability
- C. Dual power supplies
- D. Load balancing

Practice Question 1 Explanation

RAID 5, also known as disk striping with parity, uses three or more disks and writes data across all of those disks but includes additional elements known as parity blocks spread across the disks. If one of the disks fails, the system can regenerate its contents by using that parity information. Therefore, the minimum number of disks required is three.

It is not possible to use one or two disks with RAID 5 because at least three disks are required to distribute both data and parity across the array.

It is possible to use RAID 5 with five disks, but that is more than the minimum necessary to support the technology.

Correct Answer: C. 3

Practice Question 2 Explanation

Load balancing uses multiple systems to spread the burden of providing service across those systems, providing a scalable computing environment. Kevin is adding multiple servers that will share the load, so he is achieving load balancing.

If Kevin were to add more servers than necessary, that would be high availability. For example, if there were four servers, the environment would still be able to meet the full demand even if one of those servers failed. Since Kevin is only adding the three servers required to meet the load, this situation does not meet the definition of high availability.

Fault tolerance is when you add redundant components to a single server to reduce the likelihood that it will fail. Dual power supplies are an example of fault tolerance. The question does not mention dual power supplies or any other fault tolerance technology.

Correct Answer: D. Load balancing

CCSM *Certified in Cybersecurity Study Guide*, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 10 Disaster Recovery: Objective 2.2

Understand Redundancy

Business continuity programs are designed to keep a business up and running in the face of an incident or a disaster. Unfortunately, however, they may not be effective in certain circumstances. Sometimes continuity controls fail or the sheer magnitude of a disaster overwhelms the organization's capacity to continue operations. That's where disaster recovery begins.

In this chapter, you'll learn about one of the subobjectives of CC objective 2.2. The remaining material for this objective is covered in [Chapter 9, "Business Continuity."](#) The following subobjective is covered in this chapter:

Disaster Recovery (DR)

Disaster Recovery Planning

Disaster recovery (DR) is a subset of business continuity activities designed to restore a business to normal operations as quickly as possible following a disruption. Disaster recovery plans (DRPs) may include immediate measures to get operations back up and running temporarily, but DR efforts are not complete until the organization is fully back to normal.

Threat Sources

The DRP can be triggered by an environmental natural disaster, such as a hurricane; a technological failure, such as a power outage; a health emergency, such as a pandemic; or a human-caused hazard, such as a ransomware attack.

The source of a threat may be internal to the organization, such as a data center failure, or external, such as a supply chain disruption. In any case, the organization must quickly recognize the circumstances and activate its corresponding DRP.

Initial Response

Once a DRP is activated, the initial response following an emergency disruption is designed to contain the damage to the organization and to restore whatever capacity can be restored immediately.

The activities during this initial response will vary widely depending on the nature of the incident. They may include activating an alternate processing facility, containing physical damage, and calling in contractors to begin an emergency response.

Staffing

During a DR effort, the organization's focus shifts from normal business activity to a concentrated effort to restore operations as quickly as possible. From a staffing perspective, this means that many employees will be working in temporary jobs that may be completely different from their normally assigned duties. Flexibility is key during a disaster response. Also, the organization should plan disaster responsibilities as far in advance as possible and provide employees with training to prepare them to do their part during DR.

Communication

Communication is crucial to DR efforts. Responders must have secure, reliable means to communicate with each other and the organization's leadership. This includes the initial communication required to activate the DR process, even if the disaster occurs after hours; regular status updates for both field employees and leadership; and ad hoc communications to meet tactical needs.

Assessment

After the immediate danger to the organization has passed, the DR team shifts from initial response to assessment mode. The goal of this phase is simple—to triage the damage to the organization and implement functional recovery plans to restore operations continuously. In some circumstances, it may also include intermediate steps that temporarily restore operations on the way to sustainable, stable recovery.

Disaster Recovery Metrics

The following three metrics are used to help an organization plan DR efforts:

The *recovery time objective (RTO)* is the target time to restore a service to operation following a disruption.

The organization must also consider the amount of data it needs to restore. The *recovery point objective (RPO)* is the maximum time window during which data may be lost due to a disaster.

The *recovery service level (RSL)* is the percentage of a service that must be available during a disaster. For example, you might set the RSL for your website at 50%, recognizing that diminished capacity is acceptable during a disaster response.

Together, the RTO, RPO, and RSL provide valuable information to DR planners.

Exam Tip

Remember, DR efforts conclude only when the organization is back to normal operations in its primary operating environment.

Training and Awareness

Training and awareness efforts are critical components of a DRP. All personnel involved in DR efforts should receive periodic training about their role in the plan. They should also participate in more frequent awareness programs to keep their DR responsibilities top of mind.

Backups

Backups are perhaps the most important component of any DRP because most businesses today are built around their data. Whether it's proprietary designs, confidential customer lists, or information databases, data drives businesses.

For many organizations, the complete loss of data would be a disaster of tremendous proportions.

Backups provide organizations with a fail-safe way to recover their data in the event of a system failure, human error, natural disaster, or other circumstances that result in its accidental or intentional deletion or modification. Backups are a crucial safety net for data-driven businesses.

Backup Media

Organizations can back up their data in many ways. The simplest approach is to copy files from one location to another, but it is manual and error-prone. Most organizations use a more sophisticated backup strategy.

Tape Backups

Traditionally, organizations wrote their backups to magnetic tapes, and this is still a very common practice today. Linear Tape-Open (LTO) tapes, such as those shown in [Figure 10.1](#), are commonly used for this purpose. However, tapes are difficult to manage, and modern backup approaches often use alternative storage that has become much less expensive over the past few years.



Figure 10.1 LTO backup tapes.

Disk Backups

Some organizations perform *disk-to-disk backups*, which write data from the primary disk to dedicated backup disks. These backup disks may be in a separate facility, where it is unlikely that the same physical disaster would affect both the primary and backup sites. Backups sent to a storage area network (SAN) or network-attached storage (NAS) also fall into this category.

Cloud Backups

A newer trend in backups is to write backups directly to storage provided by cloud computing vendors, such as Amazon Web Services (AWS), Microsoft Azure, or their competitors. Cloud-based storage provides great geographic diversity, as the backup data is stored in separately managed facilities, and cloud providers usually perform their own backups of their systems, providing an added layer of protection for customer data.

Backup Types

There are three primary backup types—full, differential, and incremental. They differ based on the data they include.

Full Backups

Full backups, as the name implies, include everything on the media being backed up. They make a complete copy of the data.

Snapshots are a form of full backup created using specialized hardware functionality. For example, virtualization systems often provide snapshot capabilities that allow administrators to quickly create a backup disk image.

Differential and Incremental Backups

Differential backups supplement full backups and create a copy of only the data that has changed since the last full backup. *Incremental backups* are similar to differential backups, but with a small twist: They include only those files that have changed since the most recent full or incremental backup.

Backup Scenario

Let's take a quick look at an example. Joe, the storage administrator for his company, performs a full system backup every Sunday afternoon. He then performs differential backups every weekday evening. If the system fails on Friday morning, which backups would he need to restore?

First, Joe needs a base, so he would need to restore the most recent full backup from Sunday evening.

Next, Joe needs to retrieve the data that has changed since Sunday. Because Joe is using differential backups, each differential backup contains all of the data changed since the last full backup, so Joe only needs to restore the most recent differential backup—the one from Thursday evening.

What if we changed this question by switching Joe's strategy from daily differential backups to daily incremental backups? Now, Joe has a different situation on his hands. Incremental backups are smaller than differential backups and contain only those files that have changed since the most recent full *or* incremental backup.

So, Joe begins the same way, by restoring Sunday's full backup. But then he must apply each incremental backup, in order, that took place since the full backup. This means that he must apply the incremental backups from Monday, Tuesday, Wednesday, and Thursday. It takes a longer time to restore from incremental backups because of this process, but the trade-off is that incremental backups consume less space than differential backups.

Disaster Recovery Sites

During a disaster, organizations may need to shift their computing functions from their primary data center to an alternate facility designed to carry the load when the primary site is unavailable or nonfunctioning. DR sites are alternate processing facilities specifically designed for this purpose.

Most of the time, DR sites sit idle, waiting to step in when an emergency situation arises. There are three main types of alternate processing facilities:

Hot sites

Cold sites

Warm sites

Hot Sites

Hot sites are the premier form of alternate processing facility. They are fully operational data centers with all the equipment and data required to run operations.

Technology staff can activate the hot site within a moment's notice, and in many cases, the hot site will activate itself if the primary site fails. This provides an unparalleled level of redundancy, but it also comes at great expense. The costs of building and maintaining a hot site are typically similar to those of running the primary data center. You're doubling your costs to achieve tremendous recovery ability.

Cold Sites

Cold sites are facilities that can be used to restore operations, but only with a significant investment of time. They're essentially empty data centers. They have the core racks, cabling, network connections, and environmental controls necessary to support data center operations, but don't have the servers or data required to restore business operations.

Cold sites are far less expensive than hot sites, but activating them can take weeks or even months.

Warm Sites

Warm sites are a compromise. They do have the hardware and software necessary to support the company's operations, but they are not kept running in parallel. The hardware costs are the same as a hot site, but much less investment of time from IT staff is required.

Activating a warm site can take hours or days, depending on the circumstances.

Offsite Storage

DR sites not only provide a facility for technology operations but also serve as an offsite storage location for business data.

Backing up business data is important, and storing those backups in a secure facility that is geographically distant from the primary facility provides added assurance that the same disaster will not damage both the primary facility and the backups. This is all part of performing a site risk assessment as you select locations. This process is known as *site resiliency*.

Backups may be physically transported to the DR site on a periodic basis, or they may be transferred digitally via site replication, leveraging features built into an organization's storage or virtual machine platform.

When planning backup storage at offsite facilities, you'll want to make a strategic choice whether these backups are kept in an online or offline format. Online backups are available for restoration at a moment's notice, but they require a significant financial investment. Offline backups may require manual intervention to restore, but they are much less expensive.

Alternate Business Processes

In addition to alternate processing facilities, organizations can incorporate alternate business processes into their DRPs. For example, an organization might move to manual processes, such as a paper-based ordering process, if the electronic order management system will remain down for an extended period. Alternate business processes allow businesses to remain flexible in the event of a disaster.

Testing Disaster Recovery Plans

DRPs are critical to ensuring the continuity of business operations. Like any security control, they should be tested to ensure they function effectively and are ready to restore business operations in the event of a disruption. Each test of a DRP has the following two goals:

- To validate that the plan functions correctly and that the technology will work in the event of a disaster
- To provide an opportunity to identify necessary updates to the plan due to technology or business process changes

There are five major types of DR tests: read-throughs, walk-throughs, simulations, parallel tests, and full interruption tests.

Read-throughs

Read-throughs are the simplest form of DR testing. They're also known as checklist reviews. In this approach, DR staff distribute copies of the current plan to all personnel involved in DR efforts and ask them to review their individual procedures. Team members then provide feedback about any updates needed to keep the plan current.

Walk-throughs

Walk-throughs go a step further by bringing everyone together around the same table to review the plan. For this reason, the walk-through is also known

as a tabletop exercise. Walk-throughs achieve the same result as read-throughs but are generally more effective because they give the team the opportunity to discuss the plan together.

Simulations

A *simulation test* raises the stakes beyond a walk-through. The team is presented with a specific disaster scenario and works through the response in real time, actually carrying out some of the recovery procedures rather than only describing them. They might activate the emergency notification chain, retrieve backup media, or stand up a recovery instance in a test environment. The simulation stops short of disrupting production: the team exercises the mechanics of recovery without failing over live operations.

Parallel Tests

Read-throughs and walk-throughs are discussion-based exercises that talk through the plan without touching recovery technology, and a simulation begins to exercise that technology in a limited, controlled way. The parallel test goes further and fully activates the DRP, including standing up an alternate cloud or physical operating environment in response to a simulated disaster. A company doesn't actually switch operations to the backup environment, but the DR environment runs in parallel with the primary site.

Full Interruption Tests

The final test, the *full interruption test*, is the most effective type of DR test; however, it is also the most disruptive to normal operations. The business simulates a disaster by shutting down the primary operating environment and attempting to operate from the DR environment. This test type will highlight any deficiencies in the plan, but it could also have an adverse effect on the business. For this reason, full interruption tests are rare.

DR testing strategies often use a combination of different test types. An organization might conduct regular read-throughs and walk-throughs of the plan and supplement them with periodic simulations and parallel tests. Each

test type brings distinct benefits and helps the organization prepare for an actual disaster.

Exam Essentials

Disaster recovery (DR) is a subset of business continuity activities designed to restore a business to normal operations as quickly as possible following a disruption.

The recovery time objective (RTO) is the target time to restore a service to operation following a disruption. The recovery point objective (RPO) is the maximum time window during which data may be lost due to a disaster. The recovery service level (RSL) is the percentage of a service that must be available during a disaster.

Full backups include everything on the media being backed up. They make a complete copy of the data. Snapshots are a form of full backup created using specialized hardware functionality. Differential backups supplement full backups and create a copy of only the data that has changed since the last full backup. Incremental backups include only those files that have changed since the most recent full or incremental backup.

Hot sites are fully operational data centers with all the equipment and data required to run operations, ready to go. Cold sites are facilities that can be used to restore operations eventually, but they are essentially empty data centers. Warm sites have the hardware and software necessary to support the company's operations, but they are not kept running in parallel.

Disaster recovery tests range from least to most disruptive. Read-throughs ask team members to review the plan on their own, and walk-throughs gather the team to discuss it. Simulations have the team carry out part of the response to a scenario, parallel tests activate the recovery environment while the primary site keeps running, and full interruption tests shut down the primary site to run from the recovery environment.

Practice Question 1

Gene recently conducted an assessment and determined that his organization can be without its main transaction database for up to two hours before unacceptable business damage occurs. Therefore, the goal of his organization is to restore service within two hours. Which metric has Gene identified?

- A. RSL
- B. RTO
- C. RPO
- D. MTTR

Practice Question 2

Which type of recovery site has some or most systems in place but lacks the data needed to take over operations?

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

Practice Question 1 Explanation

The recovery time objective (RTO) is the time-based goal for restoring service after a disruption. Gene is setting the RTO at two hours to meet his organization's business needs.

The recovery point objective (RPO) is the maximum time period from which data may be lost as a result of a disaster, not the goal for restoring service.

The recovery service level (RSL) is the percentage of a service that must be available during a disaster, not the amount of time that may be taken to restore service.

The mean time to repair (MTTR) is the average amount of time required to repair a device and is not a goal for future recovery efforts.

Correct Answer: B. RTO

Practice Question 2 Explanation

Warm sites have systems, connectivity, and power, but they do not have the up-to-the-minute data to immediately take over operations.

A hot site can immediately take over operations, whereas a cold site has the space, power, and likely connectivity but requires systems and data to be put in place before use.

Cloud sites are not one of the three common types of recovery sites.

Correct Answer: B. Warm site

Chapter 11 Security Awareness: Objective 2.3

Understand Security Awareness

Security awareness training is a crucial part of any cybersecurity strategy. It's an ongoing process that equips people to recognize and respond to threats, including social engineering, which manipulates individuals into divulging sensitive information. The training also emphasizes the importance of robust password protection. Understanding and adopting these measures can significantly improve an organization's defense against cyber threats, turning every individual into a proactive participant in maintaining security.

In this chapter, you'll learn about CC objective 2.3. The following subobjectives are covered in this chapter:

Organizational culture (e.g., importance of security, security leadership)

Concepts (e.g., social engineering, password protection, phishing)

Organizational Culture

A strong security posture begins with the organization's culture. Technical controls and security tools are only effective when the people using them understand the importance of security and are committed to maintaining it. Organizational culture shapes how employees think about security in their daily work, and building a culture that prioritizes security is one of the most important steps an organization can take to protect its assets.

Importance of Security

Every employee in an organization plays a role in maintaining security. A single careless action, such as clicking on a malicious link or sharing a password, can compromise even the most sophisticated technical defenses. When security is embedded in the organizational culture, employees naturally consider the security implications of their actions and make better decisions.

Organizations that treat security as a core value, rather than an afterthought, experience fewer security incidents and respond more effectively when incidents do occur. In these organizations, security is not viewed as an obstacle to productivity but as an essential component of doing business responsibly. Employees understand that protecting sensitive information is part of their job, regardless of their specific role.

Building a security-aware culture requires consistent reinforcement. This includes regular communication about security threats and best practices, recognition of employees who demonstrate good security behavior, and clear consequences for security policy violations. When security becomes part of the organization's identity, employees are more likely to report suspicious activity, follow security procedures, and take ownership of their role in protecting the organization.

Security Leadership

Security culture starts at the top. When executives and managers demonstrate a commitment to security, it sends a powerful message to the rest of the organization. This concept, often called "tone at the top," means that leadership actively supports security initiatives, allocates appropriate resources to security programs, and follows security policies themselves.

Security leaders, including chief information security officers (CISOs) and security managers, play a critical role in shaping the organization's security culture. They are responsible for developing security strategies, communicating security expectations, and ensuring that security programs are effective. However, security leadership is not limited to those with security titles. Every manager and team lead has a responsibility to reinforce security practices within their teams.

Effective security leaders also serve as advocates for security within the organization. They help business leaders understand the risks the organization faces, translate technical security concepts into business terms, and ensure that security considerations are included in strategic decision-making. By bridging the gap between security and business objectives, security leaders help create an environment where security is valued and supported at every level.

Social Engineering

Digital threats aren't the only issue facing information security professionals seeking to protect their organizations. Some of the most dangerous risks come from the human threat of *social engineering*. These are also some of the hardest threats to protect against.

Social engineering attacks use psychological tricks to manipulate people into performing an action or divulging sensitive information, thereby undermining the organization's security. For example, an attacker posing as a help desk technician might use social engineering to trick a staff member into revealing their password over the telephone. Social engineering attacks are the online version of running a con.

Social engineering attacks are successful because they exploit the following factors:

Authority and trust

Intimidation

Consensus and social proof

Scarcity

Urgency

Familiarity and liking

Authority and Trust

Countless psychological experiments have shown that people will listen to and defer to someone who conveys an air of *authority*. Displaying outward signs of authority, such as dressing in a suit or simply appearing distinguished, creates trust among those without such symbols.

Well-known hacker Kevin Mitnick describes an example of authority and trust in his book *The Art of Intrusion* (Wiley, 2005). He tells of a social engineer who simply walked right into a casino security center and started issuing orders. Because he issued his commands with an air of authority, the staff complied.

Intimidation

The second reason that social engineering works is *intimidation*. It's simply browbeating people into doing what you want by scaring them and threatening that something bad will happen to them and/or the organization.

A social engineer might call a help desk posing as an administrative assistant and demand that they reset the password for an executive's account. When the help desk asks to speak to the executive, the assistant might just start yelling, "Do you know how busy he is? He is going to be very angry if you don't just do this for me!" That's intimidation.

Consensus and Social Proof

The third social engineering tactic is *consensus and social proof*. When people don't know how to react in a situation, they often look to others' behavior and follow their example. It's the herd mentality. This is what happens when someone is attacked in the street, and nobody calls 911.

It's also how riots occur. Most normal people would never think of burning a car or looting a store. But once the crowd gets going and they see this behavior around them, many people join in.

Scarcity

The fourth tactic is *scarcity*: making people believe that if they don't act quickly, they will miss out. You see this each time a major consumer electronics company releases a new product. Why will people wait in line overnight just to get a new phone? Because they want to get one before the phones run out.

A social engineer might use scarcity to trick someone into allowing them to install equipment in an office. Perhaps they show up with a Wi-Fi router and say that they are upgrading the Wi-Fi in adjacent offices with a brand-new technology and have a leftover router. If the office staff would like, they can install it here. If they agree, they think they're getting early access to new technology, while the hacker is establishing a foothold on the network.

Urgency

Urgency is the fifth tactic of social engineers. With this tactic, the threat actor creates a situation where people feel pressured to act quickly because time is running out.

For example, a social engineer might show up at an office and say they are a network technician who is there to perform a critical repair and needs access to a sensitive networking closet. When staff refuse to grant access, the social engineer can say they have another appointment and can't waste time. If the staff member opens the door now, the threat actor will perform the repair; otherwise, the network will probably go down, and they'll be out of luck.

Familiarity and Liking

The final tactic is simple: *familiarity and liking*. People want to say yes to someone they like. Social engineers will use flattery, false compliments, and fake relationships to get on a target's good side and influence their activities.

The best way to protect your organization against social engineering attacks is user education. Everyone in the organization must understand that social engineers use these tactics to gain sensitive information and be watchful for outsiders attempting to use them.

Phishing

Phishing is one of the most common and effective forms of social engineering. In a phishing attack, the attacker sends a fraudulent message, typically via email, that appears to come from a legitimate source. The goal is to trick the recipient into revealing sensitive information, such as passwords or financial data, or into clicking a malicious link that installs malware on their system.

Phishing attacks have become increasingly sophisticated over time. Early phishing emails were often easy to spot because of poor grammar, generic greetings, and obvious inconsistencies. Modern phishing attacks, however, can be highly convincing, using legitimate-looking email templates, spoofed sender addresses, and personalized content that makes them difficult to distinguish from genuine messages.

Types of Phishing

There are several variations of phishing attacks, each targeting victims in different ways:

Spear phishing targets specific individuals or groups within an organization. Unlike broad phishing campaigns that cast a wide net, spear phishing attacks are carefully crafted using information about the target, such as their name, job title, and professional relationships. This personalization makes spear phishing attacks significantly more effective than generic phishing attempts.

Whaling is a form of spear phishing that specifically targets high-level executives, such as CEOs and CFOs. Because executives often have access to sensitive information and the authority to approve financial transactions, they are particularly valuable targets. Whaling attacks frequently impersonate other executives, board members, or legal authorities.

Vishing (voice phishing) uses phone calls instead of email to trick victims. The attacker may impersonate a bank, government agency, or technical support representative and use urgency or authority to pressure the victim into providing sensitive information or taking harmful action.

Smishing (SMS phishing) delivers phishing messages via text messages. These messages often contain links to malicious websites or instruct the recipient to call a fraudulent phone number. Smishing exploits the trust people tend to place in text messages and the smaller screen size of mobile devices, making it harder to verify the legitimacy of links.

Defending Against Phishing

Defending against phishing requires a combination of technical controls and user education. On the technical side, organizations deploy email filtering systems to detect and block phishing messages, web filters to prevent access to known malicious sites, and multi-factor authentication (MFA) to limit the damage if credentials are compromised.

User education is equally important. Staff should be trained to recognize the warning signs of phishing, including unexpected requests for sensitive information, messages that create a sense of urgency, mismatched sender addresses, and suspicious links or attachments. Many organizations conduct

simulated phishing exercises to test employee awareness and identify individuals who need additional training.

Password Protection

Passwords remain one of the most common authentication mechanisms in use today. Despite advances in authentication technology, most systems still rely on passwords as either a primary or secondary means of verifying a user's identity. Effective password protection is therefore critical to maintaining the security of systems and data. I discussed password protection practices in [Chapter 2, “Authentication, Authorization, and Accounting \(AAA\).”](#) Your security awareness program should reinforce those best practices.

Security Education

Security depends on individual behavior. An intentional or accidental misstep by a single user can completely undermine many security controls, exposing an organization to unacceptable levels of risk. Security training programs help protect organizations against these risks.

Security education programs include the following two important components:

Security training provides staff members with the detailed information they need to protect the organization's security. These may use a variety of delivery techniques, but the bottom-line goal is to impart knowledge. Security training takes time and attention from staff.

Security awareness is meant to remind employees of security lessons they've already learned. Unlike security training, it doesn't require time to sit down and learn new material. Instead, it uses posters, videos, email messages, and similar techniques to keep security top of mind for those who've already learned the core lessons.

[Figure 11.1](#) shows an example of a security awareness poster.




Figure 11.1 A security awareness poster.

Source: U.S. Department of Commerce

Organizations can use a variety of different methods to deliver security training. This may include traditional classroom instruction, providing dedicated information security course material, or integrating security content into existing programs, such as a new-employee orientation program delivered by human resources. Students might also use online computer-based training providers to learn about information security, or attend classes offered by vendors. Whatever methods an organization uses, the goal is to impart security knowledge that employees can apply on the job.

Exam Tip

Two very important topics to include in security education programs are defending against social engineering attacks and protecting passwords from unauthorized disclosure and use. In fact, both topics are specifically mentioned in the Certified in Cybersecurity exam objectives, so be sure to remember them!

While all staff members should receive some degree of security education, organizations should also customize training to meet specific role-based requirements. For example, employees handling credit card information should receive training on credit card security requirements. Human Resources team members should be trained on handling personally identifiable information. IT staffers need specialized skills to implement security controls. Training should be custom-tailored to an individual's role in the organization.

You'll also want to think about the frequency of your training efforts. You'll need to balance the time required to conduct training with the benefit of reminding users of their responsibilities. One approach used by many organizations is to conduct initial training whenever an employee or contractor joins the organization or assumes new job responsibilities, and then use annual refresher training to cover the same material and update staff on new threats and controls. Awareness efforts throughout the year then keep this material fresh and top-of-mind.

Exam Tip

The team responsible for providing security training should review materials on a regular basis to ensure that the content remains relevant. Changes in the security landscape and the organization's business may require updating the material to remain fresh and effective.

Exam Essentials

Organizational culture is a critical foundation for cybersecurity. A strong security culture starts with leadership setting the tone at the top and treating security as a core organizational value. When security is embedded in the culture, employees naturally consider security implications in their daily work and take ownership of their role in protecting the organization.

The main strategies that social engineers leverage—authority and trust, intimidation, consensus and social proof, scarcity, urgency, and familiarity—exploit human psychology to circumvent security controls. Training programs should emphasize the nature of these tactics, teaching employees to recognize and respond appropriately to potential social engineering attempts.

Security training programs serve as the frontline defense in cybersecurity, empowering employees to recognize and prevent security incidents. Training efforts must continuously adapt to the evolving landscape of cyber threats, ensuring that employees remain informed about current risks and effective mitigation strategies.

Phishing is one of the most common social engineering attacks. Phishing uses fraudulent messages to trick staff members into revealing sensitive information or clicking malicious links. Variations include spear phishing (targeted), whaling (executives), vishing (voice), and smishing (SMS). Defense requires both technical controls and user education.

Practice Question 1

As a part of your company's security education program, you are asked to update the security training material. What is the main reason for keeping the training material updated?

- A. To maintain the interest of the trainees
 - B. To ensure the training material remains relevant in an evolving threat landscape
 - C. To provide a refresher course for the employees
 - D. To cover more security topics in the training
-

Practice Question 2

A person contacts your help desk pretending to be a top executive's assistant and demands immediate assistance in resetting the executive's password. The person warns that the executive will be very upset if their request isn't fulfilled promptly. Which social engineering tactic is being used?

- A. Intimidation
 - B. Consensus and social proof
 - C. Familiarity and liking
 - D. Scarcity
-

Practice Question 1 Explanation

While maintaining interest, providing refresher courses, and covering more topics can be beneficial, the main reason for keeping training material updated is to ensure it remains relevant in an evolving threat landscape. The nature of cyber threats is constantly changing, and outdated training material might leave employees unprepared to deal with new types of security threats.

Correct Answer: B. To ensure the training material remains relevant in an evolving threat landscape

Practice Question 2 Explanation

The scenario presents a situation where a person posing as an executive's assistant leverages the potential of the executive's anger to compel the help desk to act quickly on their demand. This tactic, where a threat or fear is used to manipulate the victim into compliance, is known as intimidation. Consensus and social proof would involve claiming that others have already approved or followed the same action. Familiarity and liking would rely on building trust through friendliness, shared connections, or personal rapport. Scarcity would involve suggesting that a limited opportunity or resource will disappear unless action is taken quickly.

Correct Answer: A. Intimidation

CCSM *Certified in Cybersecurity Study Guide*, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 12 Measuring Cybersecurity Effectiveness: Objective 2.4 Measure Cybersecurity Effectiveness

Organizations invest significant resources in their cybersecurity programs, but how do they know whether those investments are working? The answer lies in measuring cybersecurity effectiveness. By defining and tracking the right metrics, security leaders can assess the strength of their security posture, identify areas that need improvement, and communicate the value of security investments to business stakeholders. In this chapter, you'll learn about the key metrics that security teams use to measure their programs, including key risk indicators (KRIs). You'll also explore the dashboards, scorecards, and reports that bring those metrics to life for decision-makers.

In this chapter, you'll learn about CC objective 2.4. The following subobjectives are covered in this chapter:

Key metrics, Key Risk Indicators (KRI)

Dashboards, score cards, reports

Cybersecurity Metrics

Cybersecurity metrics are quantifiable measurements used to assess the effectiveness of an organization's security program. They allow security teams to track performance over time, identify trends, and make data-driven decisions about where to focus resources. Without metrics, security teams are left making decisions based on intuition rather than objective data.

Effective cybersecurity metrics share important characteristics.

They should be aligned with business goals and security objectives.

They should be clearly defined so that everyone involved understands what is being measured and how.

They should be repeatable, meaning that the same measurement process produces consistent results over time.

They should be actionable, providing information that security teams can use to make specific improvements. A metric that is interesting but does not lead to action is of limited value.

Key Performance Indicators (KPIs)

Key performance indicators (KPIs) measure how well the security program is performing against its goals and objectives. KPIs are backward-looking metrics, sometimes called lagging indicators, because they report on what has already happened. They help security teams understand whether their efforts are producing the desired outcomes.

Common cybersecurity KPIs include:

Mean time to detect (MTTD) measures the average time it takes to discover a cybersecurity incident after it occurs. A shorter MTTD indicates more effective monitoring and detection capabilities.

Mean time to respond (MTTR) measures the average time between detecting an incident and acknowledging it or beginning a response. This KPI reflects the efficiency of the incident response process.

Patch compliance rate tracks the percentage of systems that have been updated with the latest security patches within a defined time window. Organizations with high patch compliance rates are significantly less likely to experience a successful attack.

Phishing click rate measures the percentage of employees who click on links in simulated phishing emails during security awareness exercises. A declining click rate over time suggests that security awareness training is effective.

Vulnerability remediation time tracks the average time from discovery to closure. This metric helps organizations understand how quickly they can address security weaknesses.

Key Risk Indicators (KRIs)

While KPIs measure past performance, *key risk indicators (KRIs)* are forward-looking metrics that signal increasing risk exposure before incidents occur. KRIs are leading indicators that give security teams advance warning,

allowing them to take preventive action rather than simply reacting after something goes wrong.

Exam Tip

Remember the key difference between KPIs and KRIs. KPIs are lagging indicators that measure past performance (“How did we do?”), while KRIs are leading indicators that predict future risk (“What might go wrong?”). Both are important, but they serve different purposes. KPIs measure a business’s internal performance against its own objectives while KRIs provide insight into the business’s risk exposure.

For example, consider the number of Internet-facing systems with unpatched critical vulnerabilities. This is a KRI because a rising number indicates an increased likelihood of a breach. The metric itself does not describe a past incident; instead, it warns that conditions are becoming more favorable for an attacker. If the number crosses a defined threshold, the security team knows it is time to prioritize patching efforts.

Common cybersecurity KRIs include:

The number of unpatched critical vulnerabilities on production systems

The percentage of employees who have not completed the required security training

The number of third-party vendors with access to sensitive systems that have not undergone a recent security assessment

The number of user accounts with excessive or unnecessary privileges

The frequency of attempted intrusions detected at the network perimeter

Each of these KRIs points to a condition that, if left unaddressed, increases the likelihood of a security incident. Organizations typically establish threshold values for each KRI. When a KRI crosses its threshold, it triggers a review or escalation to ensure that the risk is addressed before it materializes into an actual incident.

Selecting the Right Metrics

One of the biggest challenges in measuring cybersecurity effectiveness is choosing which metrics to track. Organizations should avoid the temptation to measure everything. Having too many metrics can be just as problematic as having too few, because it becomes difficult to focus on what matters most. A common best practice is to select a small set of metrics that are directly tied to business objectives and the organization's risk appetite.

Organizations should also ensure that they use a balanced mix of KPIs and KRIs. KPIs tell the story of how the security program has performed, while KRIs provide early warning about emerging risks. Together, they give security leaders a comprehensive view of both current performance and future risk.

Dashboards, Scorecards, and Reports

Collecting metrics is only useful if the results are communicated effectively to the people who need them including risk owners and stakeholders. Security teams must be able to present their findings in formats appropriate to different audiences. A technical analyst needs detailed data, while a board of directors needs a high-level summary that connects security performance to business risk. Dashboards, scorecards, and reports are the three primary tools for communicating cybersecurity metrics.

Dashboards

A *cybersecurity dashboard* is a visual display that presents key security metrics in real time or near-real time. Dashboards typically use charts, graphs, gauges, and color-coded indicators to give viewers an at-a-glance understanding of the organization's security posture. Think of a dashboard like the instrument panel in a car: it provides immediate, visual feedback on the most critical measurements.

Dashboards are particularly useful for security operations teams that need to continuously monitor the environment. A security operations center (SOC) dashboard might display the number of active alerts, the current threat level, the status of critical systems, and trends in network traffic. Because dashboards update frequently, they allow security teams to spot anomalies and respond to emerging issues quickly.

Organizations often create different dashboards tailored to specific audiences. An executive dashboard might show a handful of high-level KRIs and KPIs, while a technical dashboard used by the security team might display detailed information about vulnerability scan results, firewall logs, and incident response metrics. Each is appropriate to the audience. Executives need high-level information to make strategic decisions, while security teams need actionable technical details to improve security.

Scorecards

A *cybersecurity scorecard* provides a high-level structured assessment of the organization's security posture measured against a defined framework or set of objectives. Unlike a dashboard, which focuses on real-time operational data, a scorecard provides a periodic snapshot to evaluate how well the organization is meeting its security goals. Scorecards are typically updated on a regular schedule, such as monthly or quarterly, and allow for trend analysis.

Many organizations align their scorecards with established frameworks such as the NIST Cybersecurity Framework (CSF). A CSF-aligned scorecard might rate the organization's maturity across each of the framework's core functions: Govern, Identify, Protect, Detect, Respond, and Recover. This approach allows the organization to identify specific areas where it is strong and areas where improvement is needed.

Scorecards often use a simple rating system, such as numeric scores, letter grades, or color codes (red, yellow, green), to make it easy for nontechnical stakeholders to quickly visualize and understand the results. A well-designed scorecard also shows trends over time, allowing leadership to see whether the organization's security posture is improving, declining, or remaining stable. Some organizations also include a return on security investment (RoSI) calculation in their scorecards to demonstrate the business value of security spending.

Reports

Cybersecurity reports are documents that provide analysis of security metrics, incidents, and trends. While dashboards and scorecards are designed for quick consumption, reports offer the narrative context and analysis needed to fully

understand the security landscape. Reports typically include explanations of what the data means, root-cause analysis of significant events, and recommendations for future action.

Organizations produce several types of security reports. Operational reports summarize day-to-day security activities, including the number of incidents handled, vulnerabilities discovered, and patches applied over a given period. Compliance reports demonstrate adherence to laws, regulatory requirements, and industry standards. Executive reports provide senior leadership and the board of directors with a summary of the organization's overall risk posture and the effectiveness of the security program.

The key to effective reporting is tailoring the content and level of detail to the audience. A report for the security operations team can include granular technical details, while a report for the board should focus on business impact, risk trends, and strategic recommendations. Regardless of the audience, all reports should be clear, accurate, and timely.

Exam Tip

Know the differences among dashboards, scorecards, and reports. Dashboards provide real-time or near-real-time visual displays of operational metrics. Scorecards provide periodic assessments of security posture against a framework or set of goals. Reports offer detailed narrative analysis with context and recommendations.

Exam Essentials

Key performance indicators (KPIs) are lagging indicators that measure past security performance, such as mean time to detect, mean time to respond, and patch compliance rates.

Key risk indicators (KRIs) are leading indicators that signal increasing risk exposure before incidents occur, such as the number of unpatched critical vulnerabilities or the percentage of employees who have not completed security training.

Dashboards provide real-time visual displays of security metrics for continuous monitoring. Scorecards provide periodic assessments of security posture against frameworks or goals, often using simple rating systems such as color codes. Reports deliver detailed narrative analysis, context, and recommendations tailored to specific audiences.

Effective metrics are clearly defined, repeatable, and actionable. Organizations should select a balanced mix of KPIs and KRIs that align with business objectives, and tailor the presentation of results to the needs of different audiences, from technical security teams to executive leadership.

Practice Question 1

Your organization's security team tracks the number of Internet-facing systems with unpatched critical vulnerabilities. This metric has been steadily increasing over the past three months. What type of metric is this?

- A. Key performance indicator (KPI)
- B. Key risk indicator (KRI)
- C. Scorecard rating
- D. Compliance metric

Practice Question 2

The CISO wants to present a quarterly assessment of the organization's cybersecurity posture to the board of directors, organized around the NIST Cybersecurity Framework (CSF) core functions. Which reporting tool is most appropriate for this purpose?

- A. A real-time security operations center (SOC) dashboard
 - B. A cybersecurity scorecard
 - C. An operational incident report
 - D. A vulnerability scan export
-

Practice Question 1 Explanation

The number of unpatched critical vulnerabilities on Internet-facing systems is a key risk indicator (KRI). KRIs are forward-looking metrics that signal increasing risk exposure. An increasing number of unpatched systems indicates growing risk of a breach.

KPIs, by contrast, measure past performance against goals. A scorecard rating is a periodic assessment, not a specific metric type. A compliance metric measures adherence to regulations, not future risk.

Correct Answer: B. Key risk indicator (KRI)

Practice Question 2 Explanation

A cybersecurity scorecard is the best choice for presenting a periodic assessment organized around a framework like the NIST CSF. Scorecards are designed to benchmark an organization's posture against defined objectives and present results in a format accessible to nontechnical audiences, like a board of directors.

A SOC dashboard focuses on real-time operational monitoring, not periodic assessment. An operational incident report provides detailed incident analysis, not a framework-based cybersecurity posture assessment. A vulnerability scan export is raw technical data, not an executive-level communication tool.

Correct Answer: B. A cybersecurity scorecard

CCSM *Certified in Cybersecurity Study Guide*, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Part III Domain 3: Identity and Access Management (IAM) Concepts

[Chapter 13 Identity Life Cycle Management](#)

[Chapter 14 Logical Access Controls](#)

Identity and Access Management (IAM) Concepts is the third domain of ISC2's Certified in Cybersecurity exam. It provides the knowledge that entry-level cybersecurity professionals need to know about managing identities and controlling access to systems and data. This domain includes the following two objectives:

3.1 Understand Identity Life Cycle Management

3.2 Understand Logical Access Controls

Questions from this domain make up approximately 20% of the questions on the CC exam.

CCSM Certified in Cybersecurity Study Guide, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 13 Identity Life Cycle Management:

Objective 3.1 Understand Identity Life Cycle Management

In this chapter, you'll learn about CC objective 3.1. The following subobjectives are covered in this chapter:

Roles definition

Provision

Review

Deprovision

Frameworks and tools

Every user, device, and service that interacts with an organization's information systems has a digital identity. Managing those identities from creation through deletion is an important responsibility in cybersecurity. The identity life cycle defines the stages that every identity passes through, ensuring that users and services have the access they need to do their jobs while preventing unauthorized access that could compromise security.

In this chapter, you'll explore each stage of the identity life cycle—from defining roles and provisioning accounts to conducting access reviews and deprovisioning identities when they are no longer needed. You'll also learn about the frameworks and tools that organizations use to manage identities at scale.

The Identity Life Cycle

The identity life cycle is the series of stages a digital identity passes through from creation to permanent deletion from an organization's systems.

Understanding this life cycle is essential because every gap or delay in the process creates a potential security vulnerability. An account that is provisioned with excessive privileges, left unreviewed for months, or not

promptly deprovisioned when an employee departs can become an entry point for attackers.

The identity life cycle has four main stages, as shown in [Figure 13.1](#): roles definition, provisioning, review, and deprovisioning. Each stage builds on the one before it. Role definitions establish the blueprint for the required access. Provisioning creates accounts and grants that access. Reviews ensure that access remains appropriate over time. Deprovisioning removes access when it is no longer needed. Let's examine each stage in detail.




Figure 13.1 The identity life cycle.

Roles Definition

Before an organization can grant access to its systems, it must define the roles within the organization and determine the level of access each role requires. *Roles definition* is the foundational stage of the identity life cycle because it establishes the access patterns used during provisioning.

A role is a collection of permissions that corresponds to a job function. For example, an organization might define roles such as “help desk analyst,” “network engineer,” “financial analyst,” and “HR manager.” Each of these roles would have a different set of access permissions based on the systems and data they need to perform their duties.

The process of roles definition typically involves collaboration between IT security teams and business unit managers. Business managers understand what their staff members need to do on a daily basis, while security teams understand how to translate those needs into appropriate system permissions. This collaboration helps ensure that roles are defined according to the principle of least privilege, granting only the minimum access necessary to perform each job function.

Well-defined roles simplify the entire identity life cycle. When a new employee joins the organization, administrators can simply assign them to the appropriate role rather than configuring individual permissions from scratch.

This approach, known as *role-based access control (RBAC)*, reduces errors and ensures consistency. It also makes access reviews more efficient because reviewers can evaluate access at the role level rather than examining every individual permission.

Provisioning

Provisioning is the process of creating a new digital identity and granting it the appropriate access rights. When a new employee joins an organization, the provisioning process creates their user account, assigns them to the appropriate roles, and grants them access to the systems, applications, and data they need to begin working. Provisioning can also apply to nonhuman identities, such as service accounts for applications or device identities for Internet of Things (IoT) equipment.

Effective provisioning requires a well-documented workflow. In many organizations, provisioning begins when the human resources (HR) department processes a new hire. This triggers a request to the IT department, which creates the user's account, assigns role-based permissions, and configures access to email, file shares, and business applications.

The provisioning process should include identity verification to confirm that the person receiving the account is who they claim to be. It should also include setting initial authentication credentials and requiring the user to change their password upon first login.

Organizations should follow the principle of *least privilege* during provisioning, granting new users only the access they need for their specific role. It is tempting to copy the permissions of an existing employee in a similar role, but this practice can lead to privilege creep over time as copied permissions accumulate beyond what is actually needed. Instead, provisioning should always reference the formally defined role permissions established during the roles definition stage.

Exam Tip

Remember that provisioning should always follow the principle of least privilege. New accounts should receive only the minimum permissions required for their role. Copying permissions from existing users can introduce privilege creep, where accounts accumulate unnecessary access over time.

Review

Access review is the ongoing process of examining user accounts and their associated permissions to verify that they remain appropriate. People change roles within organizations, take on new responsibilities, and move between departments. Without regular reviews, users may retain access to systems they no longer need, violating the principle of least privilege and increasing the organization's attack surface.

Organizations typically conduct access reviews on a periodic basis, such as quarterly or annually. During these reviews, managers review the access rights assigned to each team member and confirm that they remain appropriate. This process is sometimes called an access recertification or attestation because the manager is formally certifying that the access is correct. Any access that is no longer needed should be promptly revoked.

In addition to scheduled reviews, certain events should trigger an immediate access review. Job transfers, promotions, and changes in responsibilities all warrant a review of the user's access rights. When an employee moves to a new department, their old access should be removed, and new access appropriate to their new role should be granted. This process, sometimes called a role change review, prevents the accumulation of unnecessary privileges when old access is not removed during transitions.

Access reviews also serve as a detective control, helping organizations identify unauthorized access, orphaned accounts, and dormant accounts that may have been overlooked. *Orphaned accounts* are those belonging to users who have already left the organization, but whose accounts were never deprovisioned. *Dormant accounts* are accounts that have not been used for an extended period. Both account types pose security risks because attackers can exploit them undetected.

Deprovisioning

Deprovisioning is the process of removing a user's access to systems and data when that access is no longer needed. The most common trigger for deprovisioning is an employee's departure from the organization, whether through resignation, retirement, or termination. However, deprovisioning can also occur when an employee transfers to a new role and no longer needs their previous access, or when a contractor's engagement ends.

Timely deprovisioning is critical to organizational security. An account that remains active after its owner has left the organization is a significant vulnerability. Former employees who retain access could use it to steal data, sabotage systems, or cause other harm. Even if the former employee has no malicious intent, their abandoned account could be compromised by an attacker who discovers that no one is monitoring it. For this reason,

organizations should deactivate accounts immediately upon an employee's departure.

The deprovisioning process should be systematic and thorough. It typically begins by disabling the user's account rather than immediately deleting it, because the account and its associated data may need to be preserved for a period to meet legal, regulatory, or business requirements. After the retention period expires, the account and its data can be permanently deleted. In addition to disabling the account, the deprovisioning process should revoke access to all systems, recover any company-owned devices, disable remote access credentials, and remove the user from email distribution lists and collaboration tools.

Exam Tip

Deprovisioning should disable accounts immediately upon an employee's departure. Organizations typically disable rather than delete accounts at first because the data may need to be retained for legal, regulatory, or business purposes. Permanent deletion occurs after the retention period expires.

Identity Management Frameworks and Tools

Managing identities manually is feasible for small organizations, but as the number of users, systems, and applications grows, manual processes become error prone and unsustainable. Organizations rely on identity management frameworks to provide a structured approach and on specialized tools to automate the identity life cycle at scale.

Frameworks

Identity management frameworks provide organizations with a set of policies, standards, procedures, and guidelines for managing digital identities throughout their life cycle. The National Institute of Standards and Technology (NIST) publishes Special Publication 800-63, the Digital Identity

Guidelines, which provides detailed guidance on identity proofing and enrollment, authentication, and federation. This framework helps organizations determine appropriate levels of assurance for different types of transactions based on their risk.

ISO/IEC 24760 is an international standard that defines a framework for identity management. It describes the fundamental concepts of identity, including the attributes that make up an identity and the processes for managing identity information throughout its life cycle. Organizations can use this standard alongside other security frameworks, such as ISO/IEC 27001, to build a comprehensive identity management program that integrates with their broader information security management system.

Identity Management Tools

Directory services are the foundational technology for identity management. A directory service is a centralized database that stores information about users, groups, computers, and other network resources. Microsoft Active Directory (AD) is the most widely used directory service in enterprise environments. It allows administrators to create and manage user accounts, assign users to groups, and control access to network resources through group policies.

Identity and access management (IAM) platforms extend directory services with additional capabilities such as automated provisioning and deprovisioning, self-service password resets, single sign-on (SSO), and multifactor authentication (MFA). These platforms can integrate with HR systems to automatically trigger provisioning when a new employee is hired and deprovisioning when an employee leaves, reducing the delays and errors associated with manual processes.

Privileged access management (PAM) tools focus on managing and monitoring accounts with elevated privileges, such as administrator, system, and service accounts. These tools store privileged credentials in a secure vault, automatically rotate passwords, and record sessions for audit purposes. PAM is an important complement to the broader identity life cycle because privileged accounts represent the highest-value targets for attackers.

Exam Essentials

The identity life cycle consists of roles definition, provisioning, review, and deprovisioning. Role definitions establish the access patterns for each job function. Provisioning creates accounts and grants appropriate access. Review verifies that access remains appropriate over time. Deprovisioning removes access when it is no longer needed, such as when an employee leaves the organization.

Accounts that are not promptly deprovisioned when an employee departs create significant security vulnerabilities. Regular access reviews, also called access recertifications, help organizations identify excessive privileges, orphaned accounts, and dormant accounts. Both processes are essential for maintaining the principle of least privilege.

NIST SP 800-63 provides guidance on digital identity, while ISO/IEC 24760 defines an identity management framework. Key tools include directory services such as Active Directory (AD), identity and access management (IAM) platforms that automate provisioning and deprovisioning, and privileged access management (PAM) tools for securing high-privilege accounts.

Practice Question 1

A new employee joins the marketing department and needs access to the organization's content management system, social media tools, and shared file repository. Which stage of the identity life cycle determines the specific set of permissions this employee should receive?

- A. Deprovisioning
- B. Review
- C. Roles definition
- D. Provisioning

Practice Question 2

During a quarterly access review, a manager discovers that an employee who transferred from the finance department to the sales department six months ago still has access to the financial reporting system. What security principle has been violated?

- A. Separation of duties
 - B. Least privilege
 - C. Need to know
 - D. Defense in depth
-

Practice Question 1 Explanation

Roles definition is the stage of the identity life cycle that determines what access each job function requires. The marketing role would have a predefined set of permissions that includes access to the content management system, social media tools, and shared file repository.

Provisioning is the process of actually creating the account and granting that access, but the specific permissions are determined by the roles definition. Review evaluates whether existing access is still appropriate, and deprovisioning removes access.

Correct Answer: C. Roles definition

Practice Question 2 Explanation

The principle of least privilege requires that users have only the minimum access necessary to perform their current job duties. The employee who transferred to sales no longer needs access to the financial reporting system, so retaining that access violates least privilege. A role change review should have removed the old finance permissions when the employee moved to sales.

Separation of duties involves dividing critical tasks among multiple people. Need to know is related but refers specifically to information access based on job requirements. Defense in depth involves layering multiple security controls.

Correct Answer: B. Least privilege

CCSM *Certified in Cybersecurity Study Guide*, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 14 Logical Access Controls: Objective

3.2 Understand Logical Access Controls

[Chapter 2, “Authentication, Authorization, and Accounting \(AAA\),”](#) introduced the access control process and discussed authentication in detail. This chapter explores the authorization process further by examining several common access control models.

In this chapter, you’ll learn about CC objective 3.2. The following subobjectives are covered:

Principle of Least Privilege (PoLP)

Separation of Duties (SoD)

Access control models

Authorization

Authorization is the final step in the access control process. Once someone successfully authenticates to a system, authorization determines the privileges they have to access resources and information.

Least Privilege

Before discussing the different ways to implement authorization, it’s important to discuss an underlying principle: the *principle of least privilege (PoLP)*. This principle states that an individual should have only the minimum set of permissions necessary to perform their job duties.

Least privilege is important for the following two reasons:

Least privilege minimizes the potential damage from an insider attack. If an employee turns malicious, the damage they can cause will be limited by the privileges assigned to them by job role. It’s unlikely, for example, that an accountant would be able to deface the company website, because an

accountant's job responsibilities have nothing to do with updating web content.

Least privilege limits an external attacker's ability to quickly gain privileged access when compromising an employee's account. Unless they happen to compromise a system administrator's account, they will be limited by the privileges of the account they steal.

Separation of Duties

The *separation of duties (SoD)* principle states that no single person should possess two permissions that, when combined, allow them to perform a sensitive operation. Instead, those permissions should be separated and held by two different people. Account reviews and audits should inspect permissions to ensure that separation of duties is properly enforced.

One of the most common requirements for separation of duties comes in the world of accounting. Organizations typically separate the duties of adding new vendors to their accounting systems and authorizing payments to vendors. This separation prevents a single employee in the accounting department from creating a fake vendor and issuing payments to it in an attempt to embezzle funds.

When separation of duties is properly implemented, no single employee can create a fake vendor and issue payments to it. Instead, carrying out this theft would require the collusion of two employees with the required privileges. In this case, information security professionals would be responsible for configuring the access control system to prevent violations of separation of duties.

IT staff can also be the target of separation of duties controls. In the world of software development, this often surfaces as a requirement that developers cannot deploy their own code to production. Instead, they must submit their code to a testing and evaluation process, and only after it passes rigorous security and functionality checks does a separate person move the code into production.

Access Control Models

Access control systems are designed to enforce the principle of least privilege. Several different access control models are available in these systems.

Mandatory Access Control

Mandatory access control (MAC) systems are the most stringent type of access control. In MAC systems, the operating system itself restricts which permissions can be granted to users and processes on system resources. Permissions are granted by systems based on a series of security labels applied to users and the objects they want to access.

Users and system administrators themselves cannot modify permissions; therefore, MAC is rarely fully implemented on production systems outside of highly secure environments. MAC is normally implemented as a label-based mechanism where users and resources have security labels, and the operating system makes access control decisions by comparing those labels.

Discretionary Access Control

Discretionary access control (DAC) systems offer a flexible approach to authorization, allowing users to assign access permissions to other users; the owners of files, computers, and other resources have the *discretion* to configure permissions as they see fit for objects that they own. DAC systems are the most common form of access control because they provide organizations with needed flexibility.

Imagine if users in your organization didn't have the ability to assign file rights to other users as needed and IT had to be involved in every request. That would certainly make life difficult, wouldn't it?

Most organizations use DAC systems to restrict access to their data. Users who own files, printers, and other resources can grant permission to access those resources to other users. They do this by changing the permissions on the access control list (ACL) for that file, as shown in [Figure 14.1](#).




Figure 14.1 A Microsoft Windows access control list.

Role-based Access Control

Role-based access control (RBAC) systems simplify the management of authorizations. Instead of managing all permissions for an individual user, administrators create job-based roles and assign permissions to those roles. They can then assign users to roles. This is a little more work up front, but it makes life much easier down the road.

When a new user arrives, the administrator doesn't need to figure out all the explicit permissions that the user requires—the user just needs to be assigned to the appropriate roles, and all the appropriate permissions will follow. Similarly, when a group of users needs a new permission, the administrator

doesn't have to apply it to each user individually. The new permission can be assigned to the role, and all users with that role will automatically receive it.

Note

When designing an access control system, you need to select the approaches that best balance security requirements and business needs in your organization. On the one hand, if you choose a system that is not strict enough, you might unintentionally jeopardize your security. On the other hand, however, if you choose a system that is too strict, you might make it too difficult for people to get their work done.

Exam Essentials

The principle of least privilege states that users should have only the minimum set of privileges necessary to perform their job responsibilities.

The separation of duties principle says that no single person should possess two permissions that, in combination, allow them to perform a sensitive operation.

Mandatory access control (MAC) systems are the most stringent type of access control. In MAC systems, the operating system itself restricts the permissions that may be granted to users and processes on system resources.

Discretionary access control (DAC) systems allow users to assign access permissions to other users—the owners of files, computers, and other resources have the discretion to configure permissions as they see fit.

In role-based access control (RBAC) systems, administrators assign users to roles based upon their job responsibilities and assign the permissions necessary to carry out different jobs to those roles.

Practice Question 1

You are designing an access control system where each file is owned by an individual user and that user decides who can access the file. Which term best describes this access control system?

- A. MAC
 - B. RBAC
 - C. DAC
 - D. ABAC
-

Practice Question 2

A financial services company wants to ensure that no single employee can both create a new vendor account in the payment system and approve payments to that vendor. Which security principle is the company implementing?

- A. Principle of least privilege
 - B. Discretionary access control
 - C. Separation of duties
 - D. Role-based access control
-

Practice Question 1 Explanation

This is an example of a discretionary access control (DAC) system. The defining characteristic of a DAC system is that the owners of individual files can grant other users access to those files.

In a mandatory access control (MAC) system, file permissions are managed by the operating system and cannot be modified by individual users.

In a role-based access control (RBAC) system, users are assigned to roles based on their job responsibilities, and then permissions are granted to those roles.

In an attribute-based access control (ABAC) system, users are granted access to resources based on attributes assigned to their accounts. While ABAC is an incorrect answer choice used in this question, it is not specifically covered by the CC exam objectives.

Correct Answer: C. DAC

Practice Question 2 Explanation

The company is implementing the separation of duties principle by ensuring that the ability to create vendor accounts and approve payments is assigned to different employees. This prevents a single person from creating a fictitious vendor and then authorizing payments to that vendor, which is a common fraud scenario.

The principle of least privilege limits users to the minimum permissions needed for their jobs, but it does not specifically address splitting sensitive tasks between multiple people. Discretionary access control is a model where resource owners control access permissions. Role-based access control assigns permissions based on job roles, which could be used to implement separation of duties, but is not itself the principle being described.

Correct Answer: C. Separation of duties

CCSM *Certified in Cybersecurity Study Guide*, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Part IV Domain 4: Networking and Cloud Security Concepts

[Chapter 15 Network Security](#)

[Chapter 16 Network Security Architecture](#)

[Chapter 17 Cloud Security](#)

Networking and Cloud Security Concepts is the fourth domain of ISC2's Certified in Cybersecurity exam. It provides the knowledge that entry-level cybersecurity professionals need to know about securing computer networks and cloud environments. This domain includes the following three objectives:

4.1 Understand Network Security

4.2 Understand Network Security Architecture

4.3 Understand Cloud Security

Questions from this domain make up approximately 21% of the questions on the CC exam.

CCSM Certified in Cybersecurity Study Guide, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 15 Network Security: Objective 4.1

Understand Network Security

Networks form the technological backbone of modern organizations. Security professionals must understand how networks operate and the technologies used to protect them. This exam objective covers the core concepts of wired and wireless networking, and embedded systems.

In this chapter, you'll learn about three of the four subobjectives of CC objective 4.1. The firewalls subobjective is covered in [Chapter 16, “Network Security Architecture.”](#) The following subobjectives are covered in this chapter:

Concepts (e.g., Open Systems Interconnection (OSI) model, Transmission Control Protocol/Internet Protocol (TCP/IP) model, Internet Protocol version 4 (IPv4), Internet Protocol version 6 (IPv6), Virtual Private Network (VPN))

Wireless (e.g., Wi-Fi, Bluetooth)

Embedded systems (e.g., Industrial Control System (ICS)), Internet of Things (IoT)

Network Types

The networks in homes and offices are called *local area networks (LANs)*. LANs connect devices within the same building so they can communicate with each other and with servers, printers, and other devices in the office.

LANs are then connected to *wide area networks (WANs)*. WANs connect offices in different locations and also connect users to the Internet. When a LAN is connected to a WAN, it allows users to access the global Internet and communicate with anyone they'd like.

Wi-Fi networks create powerful wireless LANs that allow users to use smartphones, laptops, and other networked devices anywhere in their home or office. I'll talk more about Wi-Fi networking later in this chapter.

You've probably also used a couple of other wireless networking technologies. *Bluetooth* networks are *personal area networks (PANs)*. They're usually created on a computer or smartphone and designed for a single person. The main use of Bluetooth networks is to create wireless connections between a computer and its peripherals. Bluetooth allows users to connect wireless headsets to their phones and use their phones hands-free in their cars. The range of a Bluetooth network is about 30 feet (10 meters).

Near-field communication (NFC) technology allows extremely short-range wireless connections. For two devices to communicate via NFC, they must be no more than a few centimeters apart. NFC technology is often used for wireless payments and building access control systems.

Bluetooth and NFC networks aren't general-purpose networks that connect many computers together, but they're very useful for short-range applications.

TCP/IP Networking

Now that we've looked at the types of network connections, let's talk about how data flows on a network. To do that, we need to dive into a suite of protocols that you've probably heard about but might not be familiar with: *TCP/IP*. You might have heard people refer to TCP/IP networking because it's the protocol that powers the Internet and, basically, every LAN on the planet.

The acronym TCP/IP has two parts. The first part is the *Transmission Control Protocol (TCP)*, and the second part is the *Internet Protocol (IP)*. Each of these protocols plays a vital role in making sure that data gets from point A to point B.

IP

The first protocol that we will discuss in the TCP/IP suite is the Internet Protocol (IP). IP is responsible for routing information over the network. Even though it is referred to as the Internet Protocol, IP is used both on the Internet and on a LAN.

Each computer on an IP-based network has its own *IP address*. I'll talk more about those addresses later in this chapter. For now, just know that IP

addresses uniquely identify computers on a network and are the way that computers identify each other on an IP network.

When data travels over a TCP/IP network, it is broken into small *packets*. Each packet is a few kilobytes of data. If you're sending a large file consisting of many megabytes or gigabytes, it will be divided up into thousands of smaller packets that are sent over the network.

The reason that data is broken into packets like this is to make networks more reliable. If you tried to stuff megabytes into a single transmission and that transmission failed for some reason, you'd have to do the whole thing over again. If you break it into thousands of smaller packets, and one of those fails, you only need to retransmit that single small packet.

Larger packets would also clog up networks. Think about a city street. When many small cars are moving through the streets, traffic flows smoothly. Now imagine a mile-long freight train trying to drive down a city street. That would clog traffic for miles around, and no one else would be able to use the road until the train left the city. That would leave many unhappy drivers! For similar reasons, network traffic must be broken into pieces to avoid network congestion.

TCP/IP's job is to manage all of this work. You can just send a large file, and TCP/IP will handle the rest!

TCP

The second protocol we will discuss from the TCP/IP suite is the Transmission Control Protocol (TCP). TCP is responsible for setting up and tearing down connections between source and destination systems. TCP tracks the packets sent and controls the rate at which they are sent.

If a packet is lost or damaged along the way, TCP requests that the sender retransmit a new packet to replace the one that didn't arrive correctly. TCP is also responsible for correctly reordering packets at their destination, ensuring reliable delivery. It's the perfect protocol for activities where lost packets are unacceptable, such as web browsing, email, and file transfers.

UDP

The *User Datagram Protocol* (UDP) is an alternative to TCP. It has lower overhead but does not guarantee packet delivery from one system to another. UDP lacks TCP's reliability but operates much faster. UDP is commonly used in situations where speed is more important than ensuring that every single packet reaches its destination. UDP works well in situations where some packet loss is acceptable, such as streaming video content.

ICMP

The *Internet Control Message Protocol* (ICMP) is the housekeeping protocol of the Internet. It is part of the TCP/IP suite and is designed to allow networked devices and systems to communicate about the network's operation. For example, ICMP can be used to detect whether remote systems are live on the network, to discover the network path between two systems, and to report issues with network devices.

A lot of other components are involved in getting TCP/IP up and running on a network, but fortunately, you won't need to know about those for the CC exam.

OSI Model

Network models describe how all of these protocols fit together. The most common of these is the *Open Systems Interconnection (OSI) model*.

The OSI model describes networks as having seven different layers, as shown in [Figure 15.1](#).




Figure 15.1 The open systems interconnection (OSI) model.

The first layer, the *Physical layer*, is responsible for sending bits over the network using wires, radio waves, fiber optics, and other means.

The second layer, the *Data-Link layer*, transfers data between two nodes connected to the same physical network segment.

The third layer, the *Network layer*, expands networks to many different nodes. IP works at this layer.

The fourth layer, the *Transport layer*, establishes connections between systems and transfers data either reliably (with TCP) or as a best effort (with UDP).

The fifth layer, the *Session layer*, establishes, maintains, and terminates communication sessions between applications.

The sixth layer, the *Presentation layer*, handles data formatting and translation for transmission, including character encoding and encryption/decryption.

The seventh layer, the *Application layer*, determines how users interact with the data, using web browsers and other client applications. Protocols such as HTTP, SMTP, and SSH operate at the Application layer.

IP Addressing

Just like telephones use phone numbers and postal mail uses street addresses, the Internet needs an addressing scheme to ensure that data reaches its

intended destination. The addresses that are used by the Internet Protocol are known as *IP addresses*.

In most cases, IP addresses are written in the *dotted quad notation* of IPv4. This means that each address consists of four numbers, separated by periods. Each number can range between 0 and 255. For example, you might have a computer that uses the IP address 10.15.100.240.

Note

The dotted quad notation is used by IPv4, one of two common protocols in use today. IPv6, the next-generation addressing protocol, uses eight groups of four hexadecimal digits. For example, fae0:2660:a0a1:2efe:c84b:4c44:3467:a1ed is an IPv6 address.

While IPv4 remains widely deployed, IPv6 adoption continues to grow as the pool of IPv4 addresses shrinks. IPv6 offers several improvements beyond its expanded address space, including native IPsec support, simplified packet headers, and improved quality of service (QoS). Security professionals should be familiar with both IPv4 and IPv6 addressing, as modern networks increasingly operate in dual-stack environments that support both protocols simultaneously.

A system's IP address uniquely identifies it on a network. If the system is directly connected to the Internet, its IP address must not be used by any other system in the world, just as your mobile phone number is not used elsewhere.

Systems connected to private networks, such as the one in your home or office, may use private IP addresses that are reusable across other private networks. Your router or firewall handles translating those addresses to public IP addresses when you communicate over the Internet.


The following two IP addresses are involved in every network communication:

The *source address* indicates the system that is sending the information.

The *destination address* indicates the system that is receiving the information.


As two systems communicate, the source and destination addresses will switch, depending on who sends each packet. For example, imagine the

communication between a user with IP address 10.12.0.1 and a web server with IP address 10.51.1.2. When the user sends data to the web server, the source address is the user's IP address, and the destination address is the web server's IP address, as shown in [Figure 15.2](#).



[Figure 15.2](#) Communication from a user to a web server.

When the web server replies and sends data back to the user, the direction switches; the web server's IP address becomes the source address and the user's IP address becomes the destination address, as shown in [Figure 15.3](#).



[Figure 15.3](#) Communication from a web server to a user.

IP addresses can be assigned in two ways:

You can assign an IP address *statically*. This means that you go into the system's settings and manually specify its IP address. You are responsible for ensuring you choose a unique IP address within the range for that network.

You can assign IP addresses *dynamically* using the *Dynamic Host Configuration Protocol (DHCP)*. DHCP allows you to configure a pool of IP addresses, and it automatically assigns those addresses to systems as they join the network. Typically, servers are configured with static IP addresses, while end-user devices use dynamically assigned IP addresses.

If a system is not configured with a static IP address, it broadcasts on the network, searching for a DHCP server that can supply it with a dynamic address. If the system cannot find a DHCP server, it assigns itself an address using *Automatic Private IP Addressing (APIPA)*. APIPA addresses all begin with 169.254. If you see one of these addresses in use on your network, it likely means something has gone wrong, as the system with that address was unable to obtain a valid IP address. In all but the smallest networks, IP addresses should be assigned by a DHCP server rather than left to APIPA.

Identifying Valid IPv4 Addresses

One of the things that you may be asked to do on the exam is to identify which IP addresses are valid for a host on a network. If you see a question like this, you should approach it using the process of elimination. The following are a few simple rules that you can follow to eliminate invalid IP addresses:

No octet in an IP address should ever be larger than 255. If you see a number greater than 255 in an IP address, that is not a valid address, and you can immediately eliminate it as a possibility.

IP addresses starting with 127 are reserved for loopback addresses.

These IP addresses always reference the local system and are not valid network IP addresses. The most common loopback address is 127.0.0.1; when used in communication, it's equivalent to telling a system to talk to itself. You should never see an IP address beginning with 127.x.x.x as a host address for this reason.

The first number in an IP address should never be greater than 223.

While you can have values up to 255 in any octet, numbers higher than 223 in the first octet are reserved for special uses and shouldn't be assigned to systems. Addresses with first numbers between 224 and 239 are called *multicast addresses*. They are used to send messages to multiple systems simultaneously and should never be assigned to a single system. Addresses beginning with values between 240 and 255 are reserved for experimental use and, again, should not be found on individual systems.

Exam Tip

Be sure that you can identify valid host IP addresses when you take the CC exam. Remember the following three rules:

No value in an IPv4 address should ever be higher than 255.

The first value in an IPv4 address should never be 127.

The first value in an IPv4 address should never be higher than 223.

Domain Name System

Computers use IP addresses to communicate over the network, but those addresses are very difficult for people to remember. Just imagine if you had to memorize the IP address of every web server that you needed to access! The *Domain Name System (DNS)* allows people to use easily recognizable domain names in place of IP addresses.

DNS servers translate the names that you're more familiar with, such as www.certmike.com, into the IP addresses that computers use to communicate, such as 54.174.107.98. DNS is responsible for translating the *uniform resource locator (URL)* addresses used for websites into the IP addresses of the servers that support those sites.

Note

All of this terminology might be a little confusing. A domain name is a top-level name that may be registered by a business, organization, or individual. For example, certmike.com and isc2.org are both domain names. A URL is the address of a specific web page or other resource hosted on that domain. For example, <https://www.certmike.com/cc/> is a URL.

Every time you connect to a network, that network provides your computer with the IP address of a *DNS server* that it can use to look up IP addresses.

Then, whenever you type a website's domain name in your browser, your computer sends a request to the local DNS server for the IP address associated with that name.

If the DNS server knows the answer to your question, it simply responds to the request with the IP address, and your web browser can then connect to the website using that IP address. If the local DNS server doesn't know the answer to your question, it contacts other DNS servers to determine the correct answer and then responds to you.

DNS is a hierarchical system, and organizations that own domain names designate DNS servers as the authoritative sources of information about their domain names. When a local DNS server needs to perform a lookup, it asks a series of questions that eventually lead it to the definitive answer from the DNS server responsible for the domain.

Securing Wi-Fi Networks

One of the major responsibilities of IT professionals is to secure Wi-Fi networks to protect the wireless traffic they carry from eavesdropping attacks and to protect the network from unauthorized access. There are a few best practices you can follow to protect your Wi-Fi network.

Disabling SSID Broadcasting

Each Wi-Fi network has a name that identifies it to users. This is the name you see on your phone or laptop when choosing the Wi-Fi network you'd like to use. The technical term for this name is the service set identifier (SSID). By default, Wi-Fi networks advertise themselves to potential users by broadcasting their SSID, letting every device in the local area know the wireless network is available and accepting connections. If you don't want to advertise your network, you can disable SSID broadcasting, hiding the network from users who don't already know that it is there. [Figure 15.4](#) shows how macOS displays the SSIDs that are broadcast in an area.




Figure 15.4 SSIDs appearing on a macOS system.

Changing Default Passwords

Your wireless access point also has an administrative password that allows you to connect to the device and configure the wireless network and its security settings. The access point may have come from the manufacturer with a default password already set. It is printed in the user manual or on a label on the device itself. You should, as a matter of habit, immediately change these default passwords to strong passwords known only to you and other network administrators.

Authenticating Wi-Fi Users

You'll also need to decide which type of Wi-Fi network you'd like to run. Open networks are available to anyone who comes across them and would like to use them. Other networks use some type of user authentication to limit access.

Pre-shared Keys

Pre-shared keys (PSKs) are the simplest kind of wireless authentication and are commonly used on home Wi-Fi networks. In the PSK approach, the network uses an encryption key to control Wi-Fi access. Whenever a user wants to connect a device to the wireless network, they must enter the PSK on the device. If you've ever been in an office or public place where a Wi-Fi password is posted on the wall, that's an example of a PSK.

PSKs work fine, but they have major limitations that prevent them from being used on large networks:

Changing the network encryption key is a tremendous burden. Each time the key changes, users must reconfigure all wireless devices to use the new PSK. This might not be bad on a home network supporting a handful of users, but it is impractical in most business environments.

The use of a shared key prevents the identification of individual users and the restriction of access by user identity. For example, if a user leaves the organization, network administrators have no way to revoke that user's wireless network access short of changing the PSK on all wireless devices in the organization.

Enterprise Authentication

The more common approach to business Wi-Fi authentication is *enterprise authentication*. Under enterprise authentication, the organization runs an authentication server that verifies individual user credentials and ensures that only authorized users access the network. In this approach, instead of entering a PSK, users enter their individual username and password or provide other credentials to access the network.

Captive Portals

The last approach to wireless authentication uses captive portals, such as the one shown in [Figure 15.5](#). You may not be familiar with the term *captive portal*, but you've certainly seen it in use at hotels, airports, coffee shops, and other public locations. Captive portals provide authentication on unencrypted wireless networks.



Figure 15.5 A captive portal used for Wi-Fi authentication.

When a user connects to a network using a captive portal, they are redirected to a web page that requires them to authenticate before gaining access to the network. This authentication may be as simple as accepting the terms of service, or it may require an account password or even a credit card payment to escape the captive portal and use the Internet.

Encrypting Wi-Fi Networks

Network administrators can choose to enable encryption on Wi-Fi networks to protect communications from eavesdropping. Wi-Fi encryption is a best practice for network security. Encryption hides the true content of network

traffic from people without the decryption key. It takes an insecure communications technology—radio waves—and makes it secure.

WEP

The original approach to solving this problem was the *Wired Equivalent Privacy (WEP) protocol*. WEP was used for a long time but is now known to have serious security vulnerabilities. These issues are so significant that security professionals no longer consider WEP secure, so it should never be used on a modern network.

WPA

A newer technology called Wi-Fi Protected Access, or WPA, replaced WEP all the way back in 2003. This first version, simply called WPA, used the *Temporal Key Integrity Protocol (TKIP)* to add security that WEP lacked. TKIP changes the encryption key for each packet, preventing an attacker from discovering it by monitoring the network for a long time.

However, as is often the case with security technologies, vulnerabilities in WPA have now come to light, making it a poor choice for use on wireless networks.

WPA2

In 2004, *WPA2* was released as an upgrade to WPA. Instead of simply adding security to the old WEP standard, WPA2 uses an encryption protocol based on the *Advanced Encryption Standard (AES)* known as *CCMP*.

Security researchers have identified potential issues with WPA2, but it is still considered secure and remains widely used.

WPA3

As of July 2020, new wireless devices seeking Wi-Fi-certified designation must support the *WPA3* standard. WPA3 also supports the CCMP protocol, but it adds a new technology called *Simultaneous Authentication of Equals (SAE)*.

SAE is a secure key-exchange protocol that provides a more secure initial setup for encrypted wireless communications.

Exam Tip

To help you prepare for the CC exam, [Figure 15.6](#) summarizes all of this wireless encryption with a quick reference table.

STANDARD	SECURITY STATUS
OPEN NETWORK	INSECURE
WEP	INSECURE
WPA	INSECURE
WPA2	SECURE
WPA3	SECURE

[Figure 15.6](#) Wi-Fi standards summary.

VPN

Virtual private networks (VPNs) provide two important network security functions to IT administrators:

Site-to-site VPNs allow the secure interconnection of remote networks, such as connecting branch offices to corporate headquarters or to each other.

Remote access VPNs provide mobile workers with a mechanism to securely connect from remote locations back to the organization's network.

VPNs work by using encryption to create a virtual tunnel between two systems over the Internet. Everything that enters one end of the tunnel is encrypted and decrypted when it exits the other end. The systems on either

end of the tunnel know how to encrypt and decrypt the traffic, but anyone who intercepts the traffic in between has no way to read it.

VPNs require an endpoint on the remote network that accepts VPN connections. Many different devices may serve as VPN endpoints, including firewalls, routers, servers, and dedicated VPN concentrators. Organizations with high VPN traffic often choose a dedicated VPN concentrator because it efficiently handles VPN connections without burdening other network devices.

Embedded Systems

Embedded systems are special-purpose computers that are built into other devices to perform dedicated functions. Unlike general-purpose computers, embedded systems typically run specialized software and have limited user interfaces. They are found throughout modern organizations under the category of *industrial control systems*.

Industrial control systems (ICS) are specialized computing systems used to monitor and control physical processes in industrial environments such as manufacturing plants, power generation facilities, water treatment systems, and transportation networks. There are several types of ICS:

Supervisory control and data acquisition (SCADA) systems manage large-scale industrial processes that span wide geographic areas, such as electrical grids and pipeline systems. SCADA systems collect data from remote sensors and allow operators to monitor and control operations from a central location.

Distributed control systems (DCS) are used in industrial environments where the control functions are distributed across multiple controllers located near the equipment they manage, such as in oil refineries and chemical plants.

Programmable logic controllers (PLCs) are specialized industrial computers designed to control specific manufacturing processes, such as assembly lines and robotic systems.

ICS environments face unique security challenges. Many of these systems were originally designed for isolated networks and lack modern security features. They often run legacy software that cannot be easily patched, and downtime for security updates may be unacceptable in critical infrastructure environments.

IoT

The *Internet of Things (IoT)* refers to the growing collection of everyday objects that are embedded with computing capabilities and network connectivity. These smart devices range from consumer products such as home thermostats, security cameras, and voice assistants to commercial devices including smart lighting systems, connected medical devices, and building automation controllers.

All of these IoT devices come with security challenges. First, it is often difficult to update the software on these devices. While the devices might run slimmed-down versions of traditional operating systems, they often lack a way for the consumer to apply security patches. Second, these devices connect to the same networks used for other business purposes. If a smart device is compromised, it can serve as a gateway for attacks on other systems on the network. Finally, smart devices often connect back to cloud services for command and control, creating a potential pathway onto the network for external attackers.

One of the most effective ways to manage the security risks of embedded systems is to place them on isolated network segments. By placing IoT and ICS devices in their own network segments, separated from the rest of the organization's network by firewalls, administrators can limit the potential damage from a compromised device. This approach allows embedded systems to communicate with one another and the Internet as needed while strictly controlling access to and from the corporate network.

Network segmentation for embedded systems works much like the demilitarized zone (DMZ) concept used to protect servers that must be accessible from the Internet. The key principle is to limit the blast radius of any security incident involving these devices.

Exam Essentials

The Open Systems Interconnection (OSI) model describes network communications using seven layers. They are, in order, the Physical layer, Data-Link layer, Network layer, Transport layer, Session layer, Presentation layer, and Application layer.

Modern networking uses a suite of protocols called TCP/IP. This suite uses the IP protocol for routing information over networks at the Network layer. Each communication also uses a Transport layer protocol, typically TCP or UDP.

A system's IP address uniquely identifies it on a network. IPv4 addresses use the dotted quad notation of four integers ranging from 0 to 255, separated by periods. IPv6 addresses use eight groups of four hexadecimal digits separated by colons.

Wi-Fi networks use a variety of authentication technologies, including pre-shared keys (PSKs), enterprise authentication, and captive portals. In addition, Wi-Fi networks should use either WPA2 or WPA3 encryption to provide secure communications.

VPNs provide secure connectivity by creating encrypted tunnels over the Internet. Site-to-site VPNs connect remote networks to each other, while remote access VPNs allow individual users to securely connect to an organization's network from remote locations.

Embedded systems, including industrial control systems (ICS), present unique security challenges. ICS environments include supervisory control and data acquisition (SCADA) systems, distributed control systems, and programmable logic controllers.

ICS and IoT devices face challenges, including difficulty applying patches, shared network access, and cloud-based command-and-control. Network segmentation is a critical control for protecting embedded systems.

Practice Question 1

You are assigning a host address to a new system on a network using static IP addressing. Which one of the following is a valid IP address?

- A. 12.274.16.4
- B. 127.19.6.200
- C. 194.243.129.144
- D. 240.1.15.2

Practice Question 2

Tom is configuring a Wi-Fi network for a user who will be working from home. He would like to configure the network so that it is as strongly protected from eavesdropping as possible. Which encryption standard should Tom use?

- A. WEP
- B. WPA2
- C. WPA3
- D. WPA4

Practice Question 1 Explanation

You can walk through the three rules of valid IP addresses and answer this question by process of elimination. If you rule out any invalid IP addresses, you are left with the one valid address. First, no value in an IP address should ever be higher than 255. That eliminates the first option because it contains the value 274.

Second, the first value in a host address should never be 127. That eliminates the second option.

And finally, the first value in a host address should never exceed 223. That eliminates the last option, which contains the 240 octet value. That leaves one valid address: 194.243.129.144.

You should be prepared to answer questions like this one when you take the CC exam. If you see a question like this, be sure to read it carefully. Depending on how the question is phrased, it might ask you to identify which address in a set is valid or which is invalid. Don't get tricked by assuming what the question is asking before you read the entire thing.

Correct Answer: C. 194.243.129.144

Practice Question 2 Explanation

Tom should use WPA3 because it is the strongest available Wi-Fi encryption standard.

The WPA2 protocol is also acceptable for modern networks, but it is not as strong as WPA3, making WPA3 a better choice.

The WEP protocol is outdated and no longer considered secure; it should not be used.

The WPA4 protocol does not yet exist.

Correct Answer: C. WPA3

CCSM Certified in Cybersecurity Study Guide, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 16 Network Security Architecture:

Objective 4.1 Understand Network Security

Objective 4.2 Understand Network Security Architecture

Network security architecture encompasses the design principles and technologies that organizations use to protect their networks. This includes establishing multiple layers of defense, implementing a Zero Trust approach to access control, and dividing networks into secure segments using firewalls, VLANs, and microsegmentation.

In this chapter, you'll learn about all of the subobjectives of CC objective 4.2, as well as the firewalls subobjective of CC objective 4.1. The following subobjectives are covered in this chapter:

Firewalls (e.g., ports, applications)

Comprehending network segmentation (e.g., Firewall zones, Virtual Local Area Network (VLAN), micro-segmentation)

Defense in Depth

Zero Trust (ZT)

Defense in Depth

Defense in depth is a security strategy that uses multiple layers of security controls throughout an information technology system. The underlying principle is that no single security control is perfect, and a layered approach ensures that if one control fails, others remain in place to protect the organization's assets.

Think of defense in depth like the security of a medieval castle. The castle doesn't rely on a single wall to protect its inhabitants. Instead, it has a moat, outer and inner walls, guards at the gates, and a fortified keep at its center. An attacker who breaches one layer of defense must still contend with all the others.

In modern cybersecurity, defense in depth applies controls at multiple levels:

Physical controls include locked doors, security cameras, and badge readers that restrict physical access to facilities and equipment.

Perimeter controls include firewalls, intrusion detection and prevention systems, network gateways, and proxy servers that inspect and filter network traffic.

Network controls include network segmentation and VLANs that limit and monitor network traffic.

Endpoint controls include antimalware software, endpoint detection and response (EDR), host-based firewalls, and operating system hardening that protect individual devices.

Application controls include input validation, authentication mechanisms, and patching that protect software.

Data controls include encryption, access controls, and backups that protect data.

Administrative controls include security policies, employee training, and incident response procedures that guide human behavior.

By implementing controls at each of these layers, organizations create a comprehensive security posture that is resilient to attacks targeting any single point of failure.

Zero Trust

Traditional network security models assume that everything inside the network perimeter is trusted and everything outside is untrusted. *Zero Trust (ZT)* is a security philosophy that challenges this assumption by requiring that every user, device, and network flow be authenticated and authorized before being granted access to resources, regardless of location.

The core principle of Zero Trust is “never trust, always verify.” Under this model, even a user sitting at a workstation inside the corporate office is not automatically trusted. Every access request must be verified against security policies before it is granted.

Zero Trust architectures typically incorporate several key elements:

Strong dynamic identity verification ensures that users and devices are who they claim to be through multifactor authentication and device health checks.

Endpoint security ensures that devices meet operating system, security, and patching standards.

Least privilege access grants users and devices only the minimum permissions necessary to perform their functions, reducing the potential impact of a compromise.

Microsegmentation divides the network into small, isolated zones to contain potential breaches and limit lateral movement by attackers.

Secure communications include protection for confidentiality and integrity.

Continuous monitoring and validation ensure that trust is not granted permanently but is constantly reassessed based on the current security context.

Policy enforcement includes centralized, dynamic access policies.

Zero Trust represents a fundamental shift in how organizations approach network security. Rather than relying on a strong perimeter to keep threats out, Zero Trust assumes that threats may already be present inside the network and designs controls accordingly.

Network Segmentation

Network segmentation is the practice of dividing a computer network into smaller, isolated segments. By separating network resources into distinct zones, organizations can limit the impact of a security breach, control traffic flow between different parts of the network, and enforce security policies more effectively.

Firewalls

Firewalls are security devices that control the flow of network traffic between different network zones. They examine each network connection attempt and decide whether to allow or block it based on a set of security rules configured by the network administrator.

Ports and Applications

IP addresses can uniquely identify each system on a network, but those systems might run many different applications. That's where network ports come in. Ports allow network traffic to be directed to the correct application on a given system.

Network ports are particular locations on a system associated with a specific application. Imagine that each computer on the network is an apartment building. Every apartment building has a street address, just like every computer has an IP address. Within the building, each apartment has its own unit number. Network ports are like those apartment unit numbers: they help direct network traffic to the correct application on the system.

Network ports are represented by numbers ranging from 0 to 65,535. Different port ranges are used in different ways.

Ports between 0 and 1,023 are called *well-known ports*. These are reserved for common applications and assigned by Internet authorities. Using well-known ports helps ensure that everyone on the network knows how to reach common services.

Ports between 1,024 and 49,151 are known as *registered ports*. Application vendors can register their applications to use these ports. For example, Microsoft SQL Server uses port 1433 and Oracle databases use port 1521.

Ports above 49,151 are set aside as *dynamic ports* that applications can use on a temporary basis.

Firewall rules use network ports and application information to make access control decisions. For example, a firewall might be configured to allow traffic on port 443 (HTTPS) while blocking traffic on port 80 (HTTP) to enforce the use of encrypted web communications. Some common ports that security professionals should know include DNS (port 53), FTP (ports 20/21), HTTP (port 80), HTTPS (port 443), RDP (port 3389), SMTP (port 25), secure SMTP (port 587), and SSH (port 22).

Firewall Zones

Well-designed networks use firewalls to group systems into network segments called security zones based on each system's security level. Typical border

firewalls have three network interfaces, as shown in [Figure 16.1](#).




Figure 16.1 A network border firewall with three interfaces.

Border firewalls use these interfaces to connect three different security zones:

One interface connects to the Internet or another untrusted network. This is the interface between the protected networks and the outside world. Generally speaking, firewalls allow many different kinds of connections out to this network when initiated by a system on more trusted networks, but they block most inbound connection attempts, allowing only those that meet the organization's security policy.


A second interface connects to the organization's intranet. This is the internal network where most systems reside. This intranet zone may be further subdivided into segments for endpoint systems, wireless networks, guest networks, data center networks, and networks that meet other business needs.

The third interface connects to the *demilitarized zone (DMZ)*, also known as the *screened subnet*. The DMZ is a network segment where you can place systems that must accept connections from the outside world, such as mail or web servers. Those systems are placed in a separate security zone because they have a higher risk of compromise. If an attacker compromises a DMZ system, the firewall still blocks them from breaching the intranet.

VLANs

Virtual LANs (VLANs) are an important network security control. VLANs enable you to logically group together related systems, regardless of where they physically exist on the network. This is done at the switch level, where each switch port can be assigned to a VLAN.

Diagrams of desired network layouts typically look like the one shown in [Figure 16.2](#), with different functional groups occupying distinct network locations. Users in the HR department all share a network separate from those in the sales and IT departments.



[Figure 16.2](#) A typical network diagram.

If your building and floor layout matched the network diagram in [Figure 16.2](#) exactly, you'd be all set. More often than not, however, users from different departments are mingled together and departments are spread across buildings. That's where VLANs come into play. You can use them to connect users who are on different parts of the network to each other and to separate them from other users who might be geographically close. [Figure 16.3](#) shows how this same network might be organized using VLANs.




Figure 16.3 A typical VLAN layout.

VLANs are used to segment networks. This separates systems into networks of similar systems, reducing security risk by limiting the ability of unrelated systems to communicate with each other.

Microsegmentation

Microsegmentation is an extreme segmentation strategy that uses very small network segments that can be modified often to meet changing business requirements. With microsegmentation, each workload or application can be placed in its own isolated segment with specific security policies tailored to its needs.

While traditional segmentation using VLANs creates broad groups of similar systems, microsegmentation provides much finer-grained control. This approach aligns well with the Zero Trust philosophy discussed earlier in this chapter, as it enforces strict access controls between individual workloads rather than trusting traffic within a broad network segment.

Exam Essentials

Defense in depth is a security strategy that uses multiple layers of controls to protect assets. By implementing physical, perimeter, network, endpoint, application, data, and administrative controls, organizations ensure that the failure of any single control does not compromise the entire system.

Zero Trust (ZT) is a security philosophy based on the principle of “never trust, always verify.” Zero Trust architectures require authentication and authorization for every access request, use least privilege access, implement microsegmentation, and continuously monitor for threats.

Firewalls control network traffic between security zones by applying security rules. Organizations typically use a three-zone architecture with Internet, intranet, and demilitarized zone (DMZ) segments. Firewall rules use ports and application information to make access control decisions.

Network segmentation divides networks into smaller, isolated zones using technologies including VLANs and microsegmentation. VLANs logically group related systems regardless of physical location, while microsegmentation provides finer-grained isolation at the workload level.

Practice Question 1

You are concerned about unauthorized access to your internal network and would like to use a network device to restrict it. Which one of the following devices would best meet this need?

- A. Firewall
- B. Router
- C. Switch
- D. Access point

Practice Question 2

You are placing a new server on your organization's network. The server will host a public website and needs to be accessed by both internal employees and the public. What is the most appropriate network zone for this server?

- A. Internet
 - B. DMZ
 - C. Intranet
 - D. Extranet
-

Practice Question 1 Explanation

Looking at this question, you should be able to immediately eliminate two of the answer choices. Switches are components of an internal network; they are not connected to external networks and would not be used to control access into a network. Wireless access points are also internal network components used to connect wireless devices to the wired network. You can eliminate both of these options as possible answer choices.

That leaves two devices that do connect to external networks: firewalls and routers. The reality is that both firewalls and routers have the potential to fulfill this requirement. They both have filtering capabilities that can block unwanted external traffic. To answer this question correctly, focus on the word *best* that appears in the final sentence. The question is not asking which device *can* meet this need; it's asking which device can *best* meet this need. While routers do have limited filtering capability, they are not designed for this purpose. Firewalls, on the other hand, are designed specifically to restrict network traffic, making a firewall the best solution to this problem.

Correct Answer: A. Firewall

Practice Question 2 Explanation

The key to answering this question correctly is realizing that the server requires public access. Any server in your organization that will be accessed by the general public should be placed in the DMZ.

Servers should never be placed in the Internet zone, as this zone is completely unprotected by the firewall.

Servers should be placed on the intranet if they will be accessed only by internal users. Since this server must also be accessed by the general public, it is not appropriate to place it on the intranet.

Servers should be placed on an extranet if they will be accessed only by internal users and vendor partners. Since this server must also be accessed by the general public, it is not appropriate to place it on the extranet.

Correct Answer: B. DMZ

CCSM Certified in Cybersecurity Study Guide, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 17 Cloud Security: Objective 4.3

Understand Cloud Security

Cloud computing is one of the most transformative developments in information technology in the past decade. Organizations around the world now embrace the cloud in their IT strategies.

In this chapter, you'll learn about the cloud security concepts covered in CC objective 4.3. The following subobjectives are covered in this chapter:

Characteristics (e.g., Broad network access, rapid elasticity, measured service, on-demand self-service, resource pooling)

Service models

Deployment models

Shared security model (e.g., roles and responsibilities)

Cloud Characteristics

Cloud computing is the delivery of computing services to users over a network. Or, more formally, the National Institute of Standards and Technology (NIST) defines cloud computing as follows:

A model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction.

Organizations choose the cloud for some or all their IT workloads for a variety of reasons.

First, the cloud offers *on-demand self-service* computing. This means that technologists can access cloud resources almost immediately when they need them to do their job. That's an incredible increase in agility for individual contributors and, by extension, the organization.

Before the era of on-demand computing, a technologist who wanted to try out a new idea might have to spec out the servers required to implement the idea,

gain funding approval, order the hardware, wait for it to arrive, physically install it, and configure an operating system before getting down to work. That might have taken weeks, whereas today the same tasks can be accomplished in the cloud in a matter of seconds. On-demand self-service computing is a true game changer.

Cloud solutions also provide *scalability*. This means that, as the demand on a service increases, customers can easily increase the capacity available to them. This can occur in two different ways:

Horizontal scaling refers to adding more servers to your pool. If you run a website that supports 2,000 concurrent users with two servers, you might add a new server every time your typical usage increases by another 1,000 users. Cloud computing makes this quite easy, as you can just replicate your existing server with a few mouse clicks.

Vertical scaling refers to increasing the capacity of your existing servers. For example, you might change the number of CPU cores or the amount of memory assigned to a server. In the physical world, this means opening up a server and adding physical hardware. In the cloud, you can just click a few buttons and add memory or compute capacity.

The cloud also offers *rapid elasticity*. Elasticity is closely related to scalability. It refers to both increasing and decreasing capacity as short-term needs fluctuate. If your website experiences a surge in activity, rapid elasticity enables the system to automatically add servers until the capacity is met, then automatically remove them when the capacity is no longer needed.

The cloud offers *broad network access*. If you have the ability to access the Internet, you can connect to public cloud solutions from wherever you are—in the office, at a coffee shop, or on the road.

Finally, the cloud offers *measured service* as one of its defining characteristics. This means that almost everything you do in the cloud is metered. Cloud providers measure the number of seconds you use a virtual server, the amount of disk space you consume, the number of function calls you make, and many other measures. This allows them to charge you for precisely the services you use—no more and no less.

The measured service model can be a little intimidating when you first encounter it, but it enables cloud consumers to manage their utilization effectively and achieve the economic benefits of the cloud.

The cloud also relies on *resource pooling*. In a cloud environment, the provider's computing resources are pooled together to serve multiple consumers using a multitenant model. Various physical and virtual resources are dynamically assigned and reassigned in response to demand.

The cloud customer generally has no control or knowledge over the exact location of the provided resources but may be able to specify location at a higher level of abstraction, such as country, state, or data center. Examples of pooled resources include storage, processing, memory, and network bandwidth.

Cloud Deployment Models

When deploying cloud services, organizations have choices: private, public, hybrid, and community cloud deployment models.

Private Cloud

In the *private cloud* approach, the organization builds and runs its own cloud infrastructure or pays another organization to do so on its behalf.

Organizations using the private cloud model want to gain the flexibility, scalability, agility, and cost-effectiveness of the cloud without sharing computing resources with other organizations.

Public Cloud

The *public cloud* uses a different approach: the *multitenancy* model. In this approach, cloud providers build massive infrastructures in their data centers and then make those resources available to all comers. The same physical hardware may run workloads for many different customers simultaneously.

Multitenancy simply means that many different customers share use of the same computing resources. The physical servers that support your workloads might be the same ones supporting your neighbor's workloads.

In an ideal world, an individual customer should never see the impact of multitenancy. Servers should appear completely independent of each other and enforce the principle of *isolation*. From a security perspective, no customer

should ever be able to see another customer's data. From a performance perspective, the actions of one customer should never affect those of another. Preserving isolation is the core, crucial security task of a cloud service provider.

Of course, sometimes this concept breaks down. If cloud customers suddenly have simultaneous resource demands that exceed the environment's total capacity, performance degrades. This causes slowdowns and outages. Preventing this situation is one of the key operational tasks of a cloud service provider, and they work hard to manage workload allocation to avoid it.

When organizations use public cloud resources, they must understand that public cloud security follows a shared responsibility model. Depending on the nature of the cloud service, the cloud provider is responsible for some security areas, while the cloud customer is responsible for others.

For example, if you purchase a cloud storage service, it's your responsibility to know what data you're sending to the service and to configure access control policies that say who can access your data. It's the provider's responsibility to provide an encryption mechanism for data under their care while it's the customer's responsibility to configure and manage the encryption. It is the provider's responsibility to correctly implement access control mechanisms and the customer's responsibility to define and enforce policies as they see fit.

Cloud Connectivity

No matter which cloud deployment model you use, you'll need some way to connect your employees, office buildings, and/or physical data centers to your cloud environment. You can do this over a standard Internet connection, but many companies add a layer of security to protect that traffic from prying eyes.

The most cost-effective way to do this is to create a virtual private network (VPN) connection between your existing systems and networks to the cloud data center. If you'd like higher bandwidth, you can choose to purchase a private direct connection from your cloud provider to your enterprise network.

Hybrid Cloud

Organizations adopting a *hybrid cloud* approach use a combination of public and private cloud computing. In this model, they can use the public cloud for some workloads, while also operating their own private cloud for others, often due to data sensitivity concerns.

Community Cloud

One additional cloud model is possible, but it is not frequently used. *Community clouds* are similar to private clouds in that they are not open to the general public, but they are shared among several or many organizations that are related to each other in a common community. For example, a group of colleges and universities might come together to create a community cloud that provides shared computing resources for faculty and students across all participating schools.

Exam Tip

When taking the CC exam, remember that no one cloud model is inherently superior to the others. Some organizations might want to use a public cloud-heavy approach to achieve greater cost savings, while others might have regulatory requirements that prohibit the use of shared tenancy computing.

Cloud Service Models

Cloud services come in a variety of different categories. This section discusses the three major service models of cloud computing that you'll need to know for the CC exam: software as a service (SaaS), infrastructure as a service (IaaS), and platform as a service (PaaS).

Software as a Service (SaaS)

In a *software as a service (SaaS)* model, a public cloud provider delivers an entire application to its customers. Customers don't need to worry about processing, storage, networking, or any of the infrastructure details of the cloud service. The vendor writes the application, configures the servers, and basically gets everything running for customers, who then simply use the service. Very often, these services are accessed through a standard web browser, so very little, if any, configuration is required on the customer's end.

Common examples of SaaS applications include email delivered by Google Workspace or Microsoft 365, and storage services that facilitate collaboration and synchronization across devices, such as Box and Dropbox. SaaS applications can also be highly specialized, such as credit card processing services and travel and expense reporting.

Infrastructure as a Service (IaaS)

Customers of *infrastructure as a service (IaaS)* vendors purchase basic computing resources from vendors and piece them together to create customized IT solutions. For example, IaaS vendors might provide compute capacity, data storage, and other basic infrastructure building blocks. The three major vendors in the IaaS space are Amazon Web Services, Microsoft Azure, and Google Cloud.

Platform as a Service (PaaS)

In the third tier of public cloud computing, *platform as a service (PaaS)*, vendors provide a platform where customers can run their own application code without worrying about server configuration. This is a middle ground between IaaS and SaaS. Users don't need to worry about managing servers, but they are still running their own code.

Exam Tip

As you prepare for the CC exam, be ready to review a scenario that describes a cloud service and then select the appropriate cloud service category or deployment model described in that scenario.

Security and the Shared Responsibility Model

Security professionals need to think about security much differently in a cloud computing model. The shared security model requires that both vendors and customers take responsibility for different elements of security.

In an IaaS approach, the vendor is responsible for managing the security of their hardware and data center. Customers configure the operating system, applications, and data, so securing those elements is primarily a customer responsibility.

In a PaaS approach, the customer is still responsible for the data and applications but doesn't directly interact with the operating system, so that responsibility shifts to the vendor.

In a SaaS approach, the vendor manages almost everything, and the only responsibility that the customer has is knowing what data is stored in the service and applying appropriate access controls.

Understanding this shared responsibility model is a critical responsibility for security professionals working in a public cloud environment.

Exam Essentials

Cloud customers can choose from four deployment models. Public cloud offerings use shared hardware for many customers in a multitenant approach. Private cloud offerings dedicate hardware to a single customer. Community cloud offerings are open only to members of a specific community. Hybrid cloud offerings mix resources from two or more deployment models.

Cloud services come in three major categories: software as a service (SaaS), platform as a service (PaaS), and infrastructure as a service (IaaS).

Know the five essential characteristics of cloud computing. The five NIST essential characteristics are on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service. These characteristics define what makes a service a true cloud computing offering.

Understand the shared security model for cloud computing. In a cloud environment, the cloud provider and the customer share responsibility for security. The division of responsibilities depends on the service model. In IaaS, the customer manages more security responsibilities, while in SaaS, the provider manages most security tasks.

Practice Question 1

Your organization is shifting a series of workloads from an on-premises data center to the cloud. You are working with a vendor that allows you to start up virtual server instances as needed, and then you configure those servers to meet your business requirements. Which cloud service category does this describe?

- A. PaaS
- B. SaaS
- C. DaaS
- D. IaaS

Practice Question 2

Your organization recently signed a contract with a major IaaS provider for services that will be provided in a multitenancy model. As part of the setup process, you are working with the provider to run a dedicated fiber connection directly from your organization's headquarters campus to the cloud provider's data center. Which cloud deployment model are you using?

- A. Private cloud
- B. Community cloud
- C. Hybrid cloud
- D. Public cloud

Practice Question 1 Explanation

This is an example of an infrastructure as a service (IaaS) cloud offering because the service provider is offering you virtual server instances (a core infrastructure component), and then you are configuring those server instances to meet the needs of your organization.

In a software as a service (SaaS) approach, the provider would sell you a complete application that runs in the cloud, and you would not configure any of the underlying infrastructure yourself.

In a platform as a service (PaaS) approach, the provider would allow you to run your own application code on a platform that they manage. You still would not manage any of the underlying infrastructure.

Finally, in a desktop as a service (DaaS) offering, the cloud provider offers end-user productivity desktops that users can use in place of personal computers. That is not the situation described in this scenario.

Correct Answer: D. IaaS

Practice Question 2 Explanation

The key to answering this question correctly is to note that the services are provided in a multitenancy environment. This means that many different customers are sharing the same hardware and that this is a public cloud service offering.

A private cloud approach would use resources dedicated only to your organization and would not be multitenant. You do have a dedicated, private fiber connection to the cloud provider, but the cloud service being provided is still a public cloud offering.

There is no indication in the scenario that you are using these services in conjunction with a private cloud, which would make it a hybrid approach, or that you are using an offering open only to customers with a common background, which would make it a community cloud model.

Correct Answer: D. Public cloud

CCSM *Certified in Cybersecurity Study Guide*, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Part V Domain 5: Security Operations and Incident Response

[Chapter 18 Encryption](#)

[Chapter 19 Data Handling](#)

[Chapter 20 Security Operations](#)

[Chapter 21 Incident Response](#)

[Chapter 22 Asset Protection](#)

[Chapter 23 Security Testing](#)

Security Operations and Incident Response is the fifth domain of ISC2's Certified in Cybersecurity exam. It provides the knowledge that entry-level cybersecurity professionals need to know about everyday cybersecurity operations and responding to security incidents. This domain includes the following five objectives:

5.1 Understand Data Security

5.2 Understand Security Operations

5.3 Understand Incident Response (IR)

5.4 Understand Asset Protection

5.5 Understand Security Testing

Questions from this domain make up approximately 17% of the questions on the CC exam.

CCSM Certified in Cybersecurity Study Guide, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 18 Encryption: Objective 5.1

Understand Data Security

Encryption protects information from prying eyes by making it unreadable to anyone who does not have the required decryption key. We use encryption to protect data that is at rest on a storage device or in transit over a network.

In this chapter, you'll learn about the encryption subobjective of CC objective 5.1. The remaining material for this objective is covered in [Chapter 19, “Data Handling.”](#) The following subobjective is covered in this chapter:

Encryption (e.g., symmetric, asymmetric, hashing, quantum resistant cryptography)

Cryptography

Encryption is one of the most important security controls for protecting the confidentiality of your data. Encryption uses mathematical operations to transform information into a format where it is unreadable by anyone other than the authorized user.

Exam Tip

The math behind encryption is pretty complex, but the good news is that you won't be responsible for any of that math on the CC exam. You just need to understand the basics of encryption and how to use it to protect your organization.

Encrypting Data

When encrypting data, you start with *plaintext*. Plaintext is just a fancy way of saying the information is in its original form; anyone can read it. [Figure 18.1](#)

shows an example of a file containing plaintext—the first paragraph from this chapter.

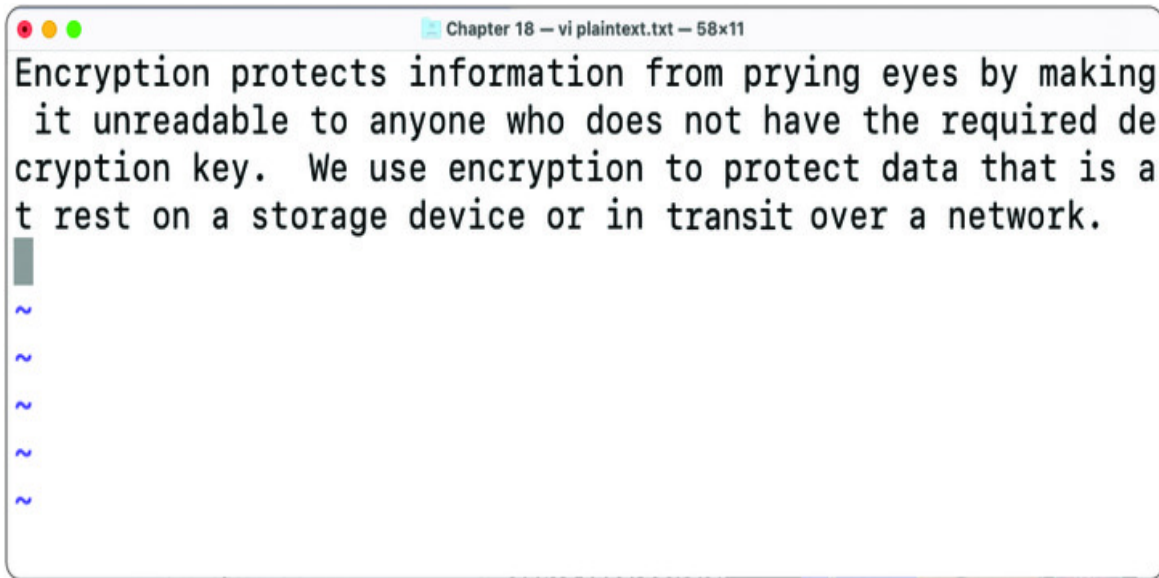


Figure 18.1 A plaintext message.

You then use an encryption algorithm, along with an encryption key, to encrypt the data. The algorithm is the mathematical formula for the encryption, and the key is basically the password to the data. Once the plaintext data has been encrypted, it is referred to as *ciphertext*. [Figure 18.2](#) shows the ciphertext obtained by encrypting the plaintext message shown in [Figure 18.1](#).

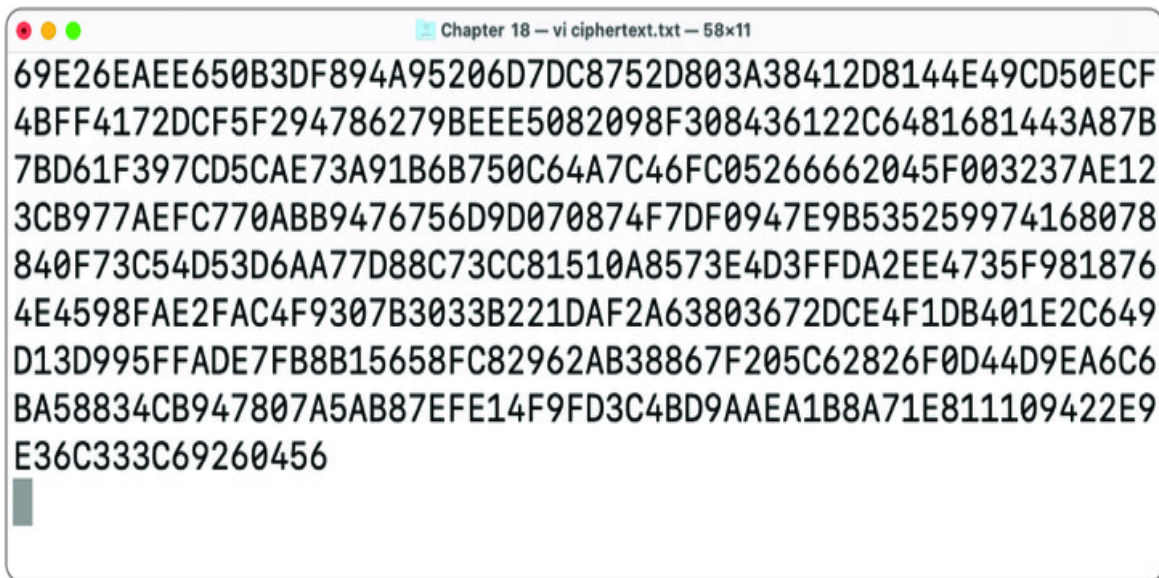


Figure 18.2 The ciphertext message obtained by encrypting the plaintext in [Figure 18.1](#).

That ciphertext data is completely unreadable to anyone who examines it. It just looks like nonsense. This stays true until the data has been decrypted.

Decrypting Data


When someone wants to *decrypt* the ciphertext and turn it back into plaintext, they use the decryption algorithm and the decryption key to perform that transformation. If they don't know the correct decryption key, the decryption simply won't work. So, protecting the decryption key is vital. As long as unauthorized users don't have access to the decryption key, they won't be able to access the data.

Encryption Algorithms

There are many kinds of encryption algorithms, and they can be categorized in different ways. Two major categories of encryption algorithms are *symmetric* and *asymmetric*.

You're probably already familiar with the concept of symmetry, meaning that two things are the same. Symmetric shapes have two sides that, when divided

along an axis, are identical. [Figure 18.3](#) shows examples of some symmetric shapes.



[Figure 18.3](#) Examples of symmetric shapes.

In cryptography, symmetry relates to keys rather than shapes.

Symmetric Encryption

In symmetric encryption algorithms, also known as shared secret encryption algorithms, encryption and decryption use the same shared secret key. If one user encrypts a message using the secret key “Apple,” a second user must decrypt it with that same secret key. It’s a shared secret. You can think of a shared secret key as the password to a message.

Let’s say Alice and Bob want to communicate. If they both know the shared secret key, they can exchange encrypted messages. This works great when only two people are involved. They can simply agree on a shared secret key and use it to encrypt messages between themselves, as shown in [Figure 18.4](#).



Figure 18.4 Symmetric encryption with two individuals.

If three people are involved, things change a little bit. Alice and Bob can still use their shared secret key to communicate privately, but now Charlie joins the picture and wants to communicate with either Alice or Bob, as shown in [Figure 18.5](#). Each person in the group wants the ability to communicate privately with any other member of the group.




Figure 18.5 Symmetric encryption with three individuals.

Alice has a way to communicate with Bob, but a second key is needed to communicate privately with Charlie. However, there is still a missing link. Bob and Charlie need a third key to communicate. For these three people to communicate, three keys are required.

As the number of participants increases, significantly more keys are needed for one-on-one communication. There's a formula that computes the number of keys required for symmetric cryptography. Where n is the number of

participants who want to communicate, the number of keys required, K , is equal to:

$$K = \frac{n(n-1)}{2}$$

As shown in [Figure 18.6](#), when more participants are involved, symmetric cryptography requires an unmanageable number of shared secret keys. For example, an organization with 10,000 employees would need almost 50 million encryption keys! If a new person were to join the organization, the organization would need to generate 10,000 new keys for that person to communicate privately with all other employees—and then distribute those 10,000 keys to every employee in the organization!




Figure 18.6 Symmetric encryption with more participants.

Asymmetric Encryption

Asymmetric cryptography solves the problem of scaling to more participants by using key pairs.

Each user gets two keys:

A *public key* that they freely distribute to anyone they want to communicate with

A *private key* that they must keep secret

In asymmetric cryptography, anything that is encrypted with one key from the pair can be decrypted with the other key from that pair. For standard communication, the sender would encrypt the message with the recipient's

publicly known public key. The recipient would then use their own private key to decrypt the message.

Exam Note

Remember that in asymmetric cryptography, the keys must be from the same pair! If Bob encrypts a message for Alice, he uses Alice's public key, and then Alice uses her own private key to decrypt the message because Alice's public and private keys come from the same pair. People get this confused on the exam all the time.

Asymmetric cryptography is slower than symmetric cryptography, but it solves the problem of key distribution for large organizations. Only two keys are needed for each user. This results in much more manageable key counts for large organizations!

Exam Note

As you prepare to answer exam questions, you should be familiar with the most common examples of symmetric and asymmetric encryption algorithms. The Advanced Encryption Standard (AES) is the most common symmetric encryption algorithm. The Rivest, Shamir, and Adleman (RSA) algorithm is the most common asymmetric encryption algorithm.

Uses of Encryption

There are two major use cases for encryption: protecting data at rest and in transit.

Data at Rest

Data at rest is data that is stored. You can encrypt individual files, folders, entire disks, or the contents of mobile devices. If someone gains access to one of those encrypted files, folders, disks, or mobile devices, you don't need to worry about the encrypted data because they won't have access to the decryption key.

Full-disk encryption (FDE) is a technology built into some operating systems that automatically encrypts all data on a device. FDE technology is particularly useful for laptops, tablets, smartphones, and other devices that might be lost or stolen. Someone who steals the device will not be able to access its data if that data is protected by full-disk encryption.

Data in Transit

Data in transit is data that's moving over a network. You can also protect this data using encryption.

When you access a website using the standard *HTTP* protocol, that data is unencrypted, and anyone who observes your network activity can eavesdrop on your web use. However, if you use the secure *HTTPS* protocol, that connection is encrypted, and the data being sent over the network is safe from prying eyes.

You can use the same encryption technology to protect the data being sent in email messages, from mobile applications to their servers, or even protect entire network connections with encryption using a *virtual private network (VPN)*.

Hash Functions

Hash functions are extremely important for public key cryptography, particularly for creating digital signatures and certificates. Let me start by giving you the technical definition of a hash function, and then I'll explain it piece by piece. A hash function is a one-way function that consistently transforms a variable-length input into a unique, fixed-length output.

Let's pick apart that definition:

Hash functions are one-way functions. That means you can't reverse the hashing process. If you have content, you can use a hash function to calculate the hash value of that content, but you can't go the other way around. If you have a hash value, you can't use it to figure out the original text.

Hash functions map variable-length input to fixed-length outputs. That simply means that you can send input of any length to a hash function and the hashes that it produces will always be the same length. Feed in two words or an entire book, and you'll get output that is the same length. That length depends on the hash function, but it will always be fixed.

Secure hash functions also produce unique output. That means that you should not be able to find two different inputs that produce the same hash value as output.

Hash functions are repeatable in that you always get the same output from hashing the same input using the same hash function.

In order for a hash function to be effective, it must meet all four of those criteria.

A hash function can fail in one of the following two ways:

First, if the hash function is reversible, it is not secure. Hash values can become public, so you don't want any way for someone to see the hash value and determine the content of a message.

The more common way that a hash function will fail is that someone will demonstrate that it is not collision-resistant. That means it doesn't meet the "unique output" part of the definition, and it is possible to find two inputs that produce the same hash. If that were the case, it would enable the forging of digital signatures and certificates. That's clearly undesirable!

Exam Note

When taking the CC exam, you should be familiar with the details of common hash functions, with particular attention to knowing which hash functions are still considered secure.

MD5

Ron Rivest created the *Message Digest 5 (MD5)* hash function in 1991. That's the same Rivest who co-invented the RSA encryption algorithm. MD5 was the fifth in a series of hash functions that grew increasingly secure. MD5 replaced MD4 after research showed MD4 was insecure.

Note

Message digest is just another term for *hash*. The two terms mean the same thing and can be used interchangeably.

MD5 produces a 128-bit hash.

Over the years, cryptanalysts discovered a series of flaws in the MD5 algorithm that eroded its collision resistance. In 2013, three cryptanalysts discovered a technique that breaks MD5's collision resistance in less than a second on a store-bought computer. Therefore, MD5 is no longer considered secure and should not be used. However, many systems still rely on MD5 for secure applications—a very bad idea!

SHA

The *Secure Hash Algorithm (SHA)* is another series of hash functions approved by the National Institute of Standards and Technology (NIST) for use in federal computing applications.

The first version of SHA, SHA-1, produces a 160-bit hash value. Cryptanalysts have discovered increasingly severe flaws in SHA-1 over the past few years and no longer consider SHA-1 secure for use.

SHA-2 replaced SHA-1 and is actually a family of six different hash functions. The SHA-2 algorithms have different hash lengths, including 224, 256, 384, and 512 bits.

All of the SHA-2 algorithms are mathematically similar to SHA-1 and MD5 and are theoretically susceptible to the same flaws that broke those algorithms.

Some attacks now exist against certain configurations of SHA-2, but the algorithm is still widely used.

NIST recognized that the mathematical similarity between SHA-2 and other, insecure, algorithms represents a future risk to SHA-2 and, thinking ahead, they began a competition to select a third version of SHA, SHA-3, in 2007. In 2012, NIST announced the selection of the Keccak algorithm as the SHA-3 standard.

SHA-3 uses a completely different mathematical technique and can actually produce a hash of any desired length. The length is set by the person computing the hash, so Keccak still satisfies the fixed-length criteria. It just allows the use of *any* fixed-length output.

Exam Tip

When taking the CC exam, you should understand that encryption is used to transform plaintext into ciphertext and that it protects both data at rest—stored in files or on devices—and data in transit as it is being sent over a network.

Quantum-Resistant Cryptography

Quantum computing is an emerging technology that uses the principles of quantum mechanics to process information in fundamentally different ways than traditional computers. While a classical computer stores data as bits that are either 0 or 1, a quantum computer uses quantum bits (qubits) that can represent both values simultaneously, allowing it to explore many possible solutions at once.

Quantum computers excel at solving certain types of mathematical problems, including those that underpin today's public-key algorithms such as RSA and elliptic curve cryptography (ECC). A sufficiently powerful quantum computer could break these algorithms, creating a long-term risk to the confidentiality and trust they provide.

Nobody knows exactly when a quantum computer capable of breaking modern cryptography will exist, and researchers still face major engineering challenges. The threat model is real, but the practical ability to execute these attacks at scale is not here today. Organizations should not panic, but they should prepare for the era of quantum computing because migrating to new cryptography takes a long time.

There is a present-day concept that makes quantum threats relevant right now, even before quantum computers can break encryption. It is called *harvest now, decrypt later (HNDL)*. In this attack strategy, an adversary captures encrypted traffic today and stores it, waiting for the day when quantum decryption becomes feasible.

When the technology catches up, the attacker can go back and read what was once protected. This is why HNDL matters most for data with a long shelf life, such as medical records, legal documents, trade secrets, and sensitive government communications.

The main response to the quantum threat is *post-quantum cryptography (PQC)*, also called *quantum-resistant cryptography*. These are new algorithms designed to resist known quantum attacks. The National Institute of Standards and Technology (NIST) has been leading the effort to standardize these algorithms and explicitly cites HNDL as a reason to move sooner rather than later. In 2024, NIST published its first three PQC standards covering key exchange and digital signatures.

Organizations should begin planning their transition to quantum-resistant cryptography now, rather than waiting for quantum computers to arrive. The migration is not a weekend project. It starts with a cryptographic inventory to identify where public-key cryptography is used across the environment, followed by prioritizing the protection of long-life secrets that are most vulnerable to HNDL attacks.

Organizations should also build crypto agility into their systems, designing architectures that can swap algorithms without rewriting everything, so they are prepared for future transitions.

Exam Essentials

Encryption is the process of transforming plaintext into ciphertext that unauthorized individuals can't read. Decryption uses a decryption key to convert the ciphertext back into plaintext.

Encryption can protect data at rest by encrypting individual files or folders. Full-disk encryption (FDE) protects the entire contents of a device by encrypting it.

Encryption can be used to protect data in transit by encrypting it as it travels over a network. The HTTPS protocol encrypts web traffic and is a secure alternative to the unencrypted HTTP protocol. Virtual private networks (VPNs) encrypt entire network connections.

Symmetric encryption algorithms, such as AES, use the same key for both encryption and decryption. Asymmetric encryption algorithms, such as RSA, use public-private key pairs.

Hash functions are one-way functions that transform a variable-length input into a unique, fixed-length output.

Quantum computing poses a future threat to public-key encryption algorithms like RSA and ECC. In the harvest now, decrypt later (HNDL) attack strategy, adversaries collect encrypted data today, intending to decrypt it with future quantum computers. Post-quantum cryptography (PQC) standards address this threat.

Practice Question 1

You are concerned that users traveling with laptops will have those devices stolen while in transit. Which encryption technology would best protect the data stored on those laptops?

- A. HTTPS
- B. HTTP
- C. VPN
- D. FDE

Practice Question 2

You have a group of remote users who need to access the corporate network from their homes. You would like to protect the information that they send back and forth to the corporate network, regardless of the application that they use. Which encryption technology would best meet this need?

- A. HTTPS
- B. HTTP
- C. VPN
- D. FDE

Practice Question 1 Explanation

This question is asking you to identify an encryption technology that will protect data at rest and, in particular, the entire contents of a laptop. Full-disk encryption (FDE) is designed for this use case. FDE encrypts an entire hard drive, rendering the contents inaccessible to an unauthorized individual who steals or finds the device.

HTTP is not an encrypted protocol. HTTPS is a secure, encrypted alternative to HTTP, but it is used to protect web traffic, which is an example of data in transit, and would not be effective in protecting a lost or stolen laptop.

Similarly, virtual private networks (VPNs) protect network communications and would not be effective in protecting a lost or stolen laptop.

Correct Answer: D. FDE

Practice Question 2 Explanation

You can begin by eliminating full-disk encryption (FDE) as an answer choice because FDE is designed to protect the contents of a disk—data at rest. It does not protect data in transit over a network.

The HTTP protocol is unencrypted, so it does not provide any protection. HTTPS is a secure, encrypted alternative to HTTP, but it is used only to protect web traffic. It would protect any web communications from the remote users but would not protect other applications.

Virtual private networks (VPNs) create secure network connections between two locations and would be an effective means of protecting the communications of these remote users.

Correct Answer: C. VPN

CCSM Certified in Cybersecurity Study Guide, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 19 Data Handling: Objective 5.1

Understand Data Security

Data security is an integral aspect of maintaining an organization's integrity and trust. Understanding how to manage data throughout its life cycle, from creation to destruction, is a critical skill for any security professional. This includes not only knowing how to securely store, use, and share data, but also how to properly archive and ultimately destroy it when it's no longer needed.

In this chapter, you'll learn about the data handling subobjective of CC objective 5.1. The remaining material for this objective is covered in [Chapter 18, "Encryption."](#) The following subobjective is covered in this chapter:

Data handling (e.g., classification, labeling, masking, and sanitization)

Data Life Cycle

The data life cycle shown in [Figure 19.1](#) is a useful way to understand the process data undergoes within an organization. It covers everything from the time data is first created until the time it is eventually destroyed. You can think of it as a way of viewing the data journey from cradle to grave.



[Figure 19.1](#) Data life cycle.

Create

In the first stage of the life cycle, the organization generates new data in either an on-premises system or the cloud. The create stage also includes

modifications to existing data.

Store

From there, the second stage of the life cycle is *store*. In this stage, the organization places the data into one or more storage systems. Again, these can be on-premises or with a cloud service provider.

Use

The next stage, *use*, is where data is actively used. Users and systems view, analyze, and process data in this stage.

Share

In the fourth stage, *share*, the data is made available to other people through one or more sharing mechanisms. This might include providing customers with a link to a file, modifying access controls so that other employees can view it, or other similar actions.

Archive

When data is no longer actively used, it moves to the fifth stage: *archive*. In this stage, data is retained in long-term storage where it is not immediately accessible but can be restored to active use if necessary.

Data archiving should follow an organization's *data retention* policy. This policy should state how long the organization will preserve different types of records and when they should be destroyed. In general, organizations should dispose of records when they are no longer necessary for a legitimate business purpose.

Destroy

In the final stage of the life cycle, *destroy*, data is eventually destroyed when it is no longer needed. This destruction should be carried out using a secure disposal method.

Data destruction must be done securely to avoid situations in which an attacker obtains paper or electronic media and then manages to reconstruct sensitive data that still exists on that media in some form.

Destroying Electronic Records

The National Institute of Standards and Technology (NIST) provides guidelines for secure media sanitization in Special Publication 800-88. It includes three different activities for sanitizing electronic media:

Clearing, the most basic sanitization technique, either writes new data to the device, overwriting sensitive data, or resets the device to factory state. Clearing is effective against most types of casual analysis.

Purging is similar to clearing but uses more advanced techniques and takes longer. Purging includes cryptographic erasure of data on disk.

Destroying is the ultimate type of data sanitization. You shred, pulverize, melt, incinerate, disintegrate, or completely destroy the media so that it is totally impossible for someone to reconstruct it. The downside of destruction, of course, is that you can't reuse the media as you would with clearing or purging.

[Figure 19.2](#) shows a flowchart to help you decide which sanitization technique to use. It is based on NIST guidelines and is widely used across government and industry. You begin the flowchart at one of three locations, depending on the classification of the information on the media, and then walk through a series of decision points based on whether you plan to reuse the media and whether it will leave your organization. The flowchart then leads you to advice on either clearing, purging, or destroying the media.




Figure 19.2 NIST storage sanitization flowchart.

Note

If you don't want to handle data sanitization and media destruction internally, third-party providers can deliver outsourced services.

Destroying Paper Records

You also should destroy paper records when they reach the end of their useful life. You have some different options at your disposal here:

Cross-cut shredding cuts hard copy documents into very small pieces that are very difficult to reassemble.

Pulping uses chemical processes to remove the ink from paper and return it to pulp form for recycling into new paper products.

Incineration burns papers, although burning paper is less environmentally friendly because it creates carbon emissions, and, unlike pulping or shredding, burned paper can't be recycled.

Exam Tip

While this process is described as a life cycle, it's important to note that the stages of the life cycle are not always followed in order and not all of them occur for every piece of data. For example, it is possible to create new data in memory, use it there, and then destroy it without ever storing it in a repository. Similarly, data might be permanently retained in active storage and never reach the archive or destroy stages. However, the life cycle is still a useful model for understanding the different stages of data life.

Data Masking

Organizations often need to use data in environments where exposing the actual sensitive values would create unnecessary risk. *Data masking* is a technique that replaces sensitive data with realistic but fictitious substitutes, allowing the data to remain useful for purposes like software testing, analytics, and training without exposing the actual sensitive information.

There are several approaches to data masking that organizations may use:

Static data masking creates a permanent copy of a dataset, replacing sensitive values with masked values. For example, an organization might create a copy of its customer database, replacing all names and Social Security numbers with fictitious values for use in a software testing environment.

Dynamic data masking applies masking rules in real time as users access data, showing different versions of the data to different users based on their permissions. For example, a customer service representative might see a credit card number displayed as ****-****-****-1234, while a billing administrator sees the full number.

Tokenization replaces sensitive data elements with unique tokens that have no exploitable meaning. The original data is securely stored in a separate token vault, and authorized systems can exchange the token for its original value when needed. This approach is commonly used in payment card processing.

Two related concepts are important to understand. *Data anonymization* permanently removes all identifying information from a dataset, making it impossible to trace the data back to a specific individual. *Pseudonymization* replaces identifying information with artificial identifiers but retains the ability to re-identify individuals using a separate mapping table. Pseudonymized data still requires security controls because re-identification is possible.

Data masking differs from encryption in an important way. Encrypted data is reversible: anyone with the correct key can decrypt it and recover the original data. Masked data, however, is generally not reversible. The original values are replaced rather than transformed, making masking appropriate when there is no need to recover the original data.

Data Classification

Organizations use *classification* to help users understand the security requirements for handling different types of information. Data classification serves as the backbone of data governance efforts designed to help organizations provide appropriate protections for sensitive data.

Data classification policies describe the security levels of information used in an organization and the process for assigning information to a particular classification level. The different security categories, or classifications, used by an organization determine the appropriate storage, handling, and access requirements for classified information.

Security classifications are assigned based on the sensitivity of the information and its criticality to the enterprise.

Classification Levels

Classification levels vary but all basically try to group information into high, medium, and low sensitivity levels and differentiate between public and private information.

For example, the U.S. federal government uses the following classification levels to safeguard government data:

Top Secret

Secret

Confidential

A business, however, might use friendlier terms to accomplish the same goal, such as:

Highly Sensitive

Sensitive

Internal

Public

Businesses use these terms to describe how they will handle sensitive proprietary and customer data.

Organizations that deal with sensitive personal information should consider not only the impact on their organization if they lose control of that data but also the impact on their customers. For this reason, they should pay particular attention to safeguarding the following:

Personally identifiable information (PII)

Financial information, including credit card and bank account numbers

Health information, especially records that are governed by HIPAA privacy and security standards

Data classification is extremely important because it is used as the basis for other data security decisions. For example, a company might require strong encryption to protect Sensitive and Highly Sensitive information at rest and in transit. This is an example of a data handling requirement.

Labeling

When an organization classifies information, it should also include *labeling* requirements that apply consistent markings to sensitive information. Using standard labeling practices ensures that users can consistently recognize sensitive information and handle it appropriately.

Exam Essentials

Data must be appropriately managed throughout its life cycle—from creation to destruction. This includes securely creating, storing, using, sharing, archiving, and ultimately destroying data when it's no longer necessary. The destruction method should be secure and adhere to standards such as the NIST guidelines for media sanitization.

Data classification is crucial for implementing appropriate security measures. Information should be categorized based on its sensitivity and criticality to the organization. Common classification levels include U.S. federal government levels like Top Secret, Secret, and Confidential, and business-friendly terms such as Highly Sensitive, Sensitive, Internal, and Public. These levels determine the appropriate storage, handling, access, and protection requirements for information.

Labeling of classified information should be standardized within an organization. This practice ensures that users can consistently recognize sensitive information and handle it correctly, thereby mitigating the risk of accidental data leaks or breaches.

Data masking protects sensitive information by replacing it with realistic but fictitious values. Static masking creates permanent masked copies, dynamic masking applies rules in real time based on user permissions, and tokenization replaces sensitive values with meaningless tokens.

Anonymization permanently removes identifying information, while pseudonymization replaces it with artificial identifiers that can be reversed.

Practice Question 1

An organization classifies its data into four categories: Public, Internal, Sensitive, and Highly Sensitive. Which of the following data would likely be classified as Highly Sensitive?

- A. Press releases about new products
 - B. Internal policy documents
 - C. Customer credit card information
 - D. Employee cafeteria menu
-

Practice Question 2

You're a security officer at a corporation that has a strict data retention policy. One of the company's servers has reached end of life and is being decommissioned. Which of the following actions is the most secure method to ensure that sensitive data on the server's magnetic hard drives (HDDs) is securely sanitized if you would like to potentially reuse the drives?

- A. Overwriting the data
 - B. Formatting the drives
 - C. Degaussing the drives
 - D. Physically destroying the drives
-

Practice Question 1 Explanation

In this question, you need to identify which type of data would most likely be classified as Highly Sensitive. Press releases are public information and would fall under the Public category. Internal policy documents are typically Internal or Sensitive, depending on the content. Employee cafeteria menus are also typically classified as Public or Internal.

The only option that would likely be classified as Highly Sensitive is customer credit card information, given its nature and the significant impact its disclosure could have on customers and the organization.

Correct Answer: C. Customer credit card information

Practice Question 2 Explanation

This question asks for the most secure method to ensure data destruction while still allowing drive reuse.

Physically destroying the drives is the most secure method of data destruction, but it eliminates the possibility of reusing the drives, which disqualifies it as a viable answer.

Formatting the drives does not necessarily erase all data and may leave remnants that can be recovered.

Degaussing is a process that demagnetizes magnetic hard drives and can be quite effective, but it also renders them unusable for future use. Note that degaussing is only effective on magnetic media and does not work on solid-state drives (SSDs), hybrid drives, or flash storage.

Overwriting data is one of the most secure methods of data destruction that still allows for hard drive reuse. It ensures that the previous data is replaced with new data, making it difficult to recover the original data.

Correct Answer: A. Overwriting the data

CCSM Certified in Cybersecurity Study Guide, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 20 Security Operations: Objective 5.2

Understand Security Operations

Security operations encompass the day-to-day activities that protect an organization's information assets. This includes monitoring systems for signs of an attack, understanding the threat landscape, and using structured frameworks to analyze and respond to security events.

In this chapter, you'll learn about the subobjectives of CC objective 5.2. The following subobjectives are covered in this chapter:

Logging and monitoring security events

Security event triage (e.g., incident use cases, prioritization, correlation)

Threat actors (e.g., types, motivations)

Cyber threat intelligence

Threat frameworks

Logging

Logs provide a treasure trove of information for security professionals, whether they're investigating an incident, troubleshooting a technical problem, or gathering evidence. When logging is configured properly, organizations can look at a specific event and achieve three important objectives:

- They can determine who caused the event. That's a characteristic known as *accountability* or identity attribution. This attribution can be to a specific person, a computer's IP address, or a geographic location.
- They can track down all other events related to the investigated event. That's a characteristic known as *traceability*.
- They can provide clear documentation of those actions. That's *auditability*.

Log Monitoring

If you're like most security professionals, you simply don't have the time to thoroughly review logs. There are simply far too many log entries generated by systems, applications, network devices, and other event sources each day, and trudging through them would be tedious, mind-numbing work.

Fortunately, computers are very good at tedious work, and most organizations now go beyond simple reporting and alerting mechanisms and apply artificial intelligence approaches to the problem of security log analysis.

Security information and event management (SIEM) systems have the following two major functions on an enterprise network:

- They act as a central, secure collection point for log entries. Administrators configure all their systems, network devices, and applications to send log records directly to the SIEM, which stores them securely, protected from unauthorized modification and available for analysis.
- They correlate log entries to detect patterns of potentially malicious activity, increasingly using artificial intelligence.

The great thing about a SIEM is that it has access to all log entries across the organization. In a hierarchical organization, network engineers might have access to firewall logs, system engineers to operating system logs, and application administrators to application logs. This siloed approach means attacks could go unnoticed if their signs are spread across multiple departments. Each administrator might see a piece of the puzzle but can't put the whole picture together.

The SIEM has all the puzzle pieces and performs *log correlation* to identify combinations of activity that could indicate a security incident.

For example, a network intrusion detection system might notice the unique signature of an attack in inbound network traffic, triggering an event within the SIEM that pulls together other information. From there, a network firewall might note an inbound connection to a web server from an unfriendly country.

The web server might report suspicious queries that include signs of a SQL injection attack. The database server might report a large query from a web application that deviates from normal patterns, and a router might report a large flow of information from the database server to the Internet.

In isolation, each of these activities may seem innocuous, but when the SIEM puts those pieces together, a pattern of suspicious activity emerges.

Security Event Triage

Security operations centers (SOCs) receive thousands of alerts each day from firewalls, intrusion detection systems, endpoint protection tools, and other security controls. *Security event triage* is the process of classifying, validating, and prioritizing these alerts to determine which require immediate investigation and which can be addressed later or dismissed.

Analysts prioritize security events based on several factors. The *severity* of the event reflects its technical impact or damage potential. The *criticality* of the affected asset is determined by the business value of the system involved. For example, an alert on a public-facing web server containing customer data warrants greater urgency than the same alert on a development workstation. Analysts also consider the potential business *impact*, including possible data loss, service disruption, or regulatory consequences.

SOC teams develop *incident use cases* that describe common attack scenarios and the appropriate response procedures. These use cases help analysts quickly recognize patterns such as unauthorized access attempts, malware infections, data exfiltration, and privilege escalation. By matching incoming alerts to known use cases, analysts can respond more efficiently and consistently.

The triage process works hand in hand with the event correlation capabilities of SIEM systems described earlier in this chapter. While SIEM systems automate the correlation of related events, human analysts apply judgment and context to determine the true nature and priority of the resulting alerts.

Threat Actors

Understanding who might attack your organization is essential to building effective defenses. *Threat actors* are individuals or groups who conduct cyberattacks with the intent to cause harm. They vary widely in their motivations, capabilities, and resources.

Nation-state actors are government-sponsored groups that conduct cyber espionage, sabotage, and intelligence gathering. They are typically well-funded, highly sophisticated, and capable of sustained, long-term campaigns

known as advanced persistent threats (APTs). Their targets often include government agencies, critical infrastructure, and strategic industries.

Organized cybercriminals are motivated primarily by financial gain. They conduct ransomware attacks, phishing campaigns, credit card theft, and other schemes designed to generate profit. Many cybercriminal groups operate using a business model, offering ransomware-as-a-service (RaaS) and other criminal tools to affiliates.

Hactivists are individuals or groups motivated by political or social causes. They use cyberattacks to promote their ideology, often through website defacement, distributed denial-of-service (DDoS) attacks, or the public release of stolen data intended to embarrass their targets.

Insider threats come from individuals within the organization, such as employees, contractors, or business partners who have legitimate access to systems and data. Insiders may act out of financial motivation, revenge, or coercion. They are particularly dangerous because they already have trusted access and can bypass many security controls.

Script kiddies are inexperienced attackers who use publicly available tools and scripts to launch attacks. They are typically motivated by curiosity or a desire for notoriety rather than financial gain or ideology. Although their individual skill level is low, the widespread availability of attack tools means they can still cause significant damage.

Cyber Threat Intelligence

Cyber threat intelligence (CTI) is the collection, analysis, and sharing of information about current and potential cyber threats, including threat actors and their tactics, techniques, and procedures (TTPs). CTI helps organizations make informed decisions about defending against and responding to threats.

Threat intelligence operates at several levels. *Strategic* intelligence provides high-level information about the threat landscape and trends, helping executives make decisions about security investments and policies.

Operational intelligence provides details on specific attack campaigns, helping defenders understand how adversaries are likely to target their organization. *Tactical* intelligence focuses on the specific technical indicators

associated with threats, such as malware signatures, malicious IP addresses, and suspicious domain names.

A key concept in threat intelligence is the *indicator of compromise (IoC)*. IoCs are forensic artifacts that suggest a system has been breached or is under attack. Common examples include malicious IP addresses, file hashes associated with known malware, suspicious domain names, and unusual network traffic patterns. Organizations feed IoCs into their security tools, such as SIEM systems and intrusion detection systems, to automatically detect known threats.

Organizations obtain threat intelligence from a variety of sources, including *open-source intelligence (OSINT)* from public sources, commercial threat intelligence services from security vendors, government agencies such as the Cybersecurity and Infrastructure Security Agency (CISA), and industry-specific Information Sharing and Analysis Centers (ISACs) that facilitate threat information sharing among organizations in the same sector.

Threat Frameworks

Threat frameworks provide structured approaches for understanding and analyzing cyberattacks. They help security teams identify attack patterns, attribute threats to specific actors, and develop effective defenses.

MITRE ATT&CK

The *MITRE ATT&CK* framework is a comprehensive knowledge base of adversary tactics and techniques based on real-world observations. It organizes attack behaviors into tactics (the adversary's goals, such as initial access, persistence, and data exfiltration) and techniques (the specific methods used to achieve those goals, such as phishing, exploitation of public-facing applications, and command-and-control communications). Security teams use ATT&CK to identify gaps in their defenses, develop detection rules, and communicate about threats using a common language.

When cybersecurity analysts describe an attack, they often use the shorthand term TTPs, which stands for tactics, techniques, and procedures.

Tactics describe the attacker's objective: the "why" behind a set of actions.

Techniques describe the general method used to achieve that objective: the “how” at a high level.

Procedures describe the specific implementation details observed in the real world, such as the exact commands, tools, or scripts an attacker used.

ATT&CK organizes these behaviors into a consistent structure that defenders can use to analyze incidents and move from raw observations, such as logs and alerts, to a clear description of what the adversary did and why it mattered.

[Figure 20.1](#) shows an example of the ATT&CK technique definition for active scanning techniques. It provides an ID number, as well as classification details such as the tactic, the platforms it applies to, procedure examples, potential mitigating controls, and detection strategies. You can find the ATT&CK website at <https://attack.mitre.org>.

Active Scanning

Sub-techniques (3) ▾

Adversaries may execute active reconnaissance scans to gather information that can be used during targeting. Active scans are those where the adversary probes victim infrastructure via network traffic, as opposed to other forms of reconnaissance that do not involve direct interaction.

Adversaries may perform different forms of active scanning depending on what information they seek to gather. These scans can also be performed in various ways, including using native features of network protocols such as ICMP.^{[1][2]} Information from these scans may reveal opportunities for other forms of reconnaissance (ex: [Search Open Websites/Domains](#) or [Search Open Technical Databases](#)), establishing operational resources (ex: [Develop Capabilities](#) or [Obtain Capabilities](#)), and/or initial access (ex: [External Remote Services](#) or [Exploit Public-Facing Application](#)).

ID: T1595

Sub-techniques: [T1595.001](#), [T1595.002](#), [T1595.003](#)

① **Tactic:** [Reconnaissance](#)

① **Platforms:** PRE

Version: 1.0

Created: 02 October 2020

Last Modified: 24 October 2025

[Version Permalink](#)

Procedure Examples

ID	Name	Description
C0030	Triton Safety Instrumented System Attack	In the Triton Safety Instrumented System Attack , TEMP:Veles engaged in network reconnaissance against targets of interest. ^[3]

Mitigations

ID	Mitigation	Description
M1056	Pre-compromise	This technique cannot be easily mitigated with preventive controls since it is based on behaviors performed outside of the scope of enterprise defenses and controls. Efforts should focus on minimizing the amount and sensitivity of data available to external parties.

Detection Strategy

ID	Name	Analytic ID	Analytic Description
DET0830	Detection of Active Scanning	AN1962	Monitor network data for uncommon data flows. Processes utilizing the network that do not normally have network communication or have never been seen before are suspicious. Monitor and analyze traffic patterns and packet inspection associated to protocol(s) that do not follow the expected protocol standards and traffic flows (e.g extraneous packets that do not belong to established flows, gratuitous or anomalous traffic patterns, anomalous syntax, or structure). Consider correlation with process monitoring and command line to detect anomalous processes execution and command line arguments associated to traffic patterns (e.g. monitor anomalies in use of files that do not normally initiate connections for respective protocol(s)).

Figure 20.1 The ATT&CK definition for active scanning.

Cyber Kill Chain

The *Cyber Kill Chain*, developed by Lockheed Martin, models cyberattacks as a sequence of seven stages, as shown in [Figure 20.2](#). The key defensive insight of the Kill Chain is that disrupting an attack at any stage prevents the adversary from reaching their ultimate goal. Early detection and intervention are far more effective than waiting to respond after the attacker has achieved their objectives.

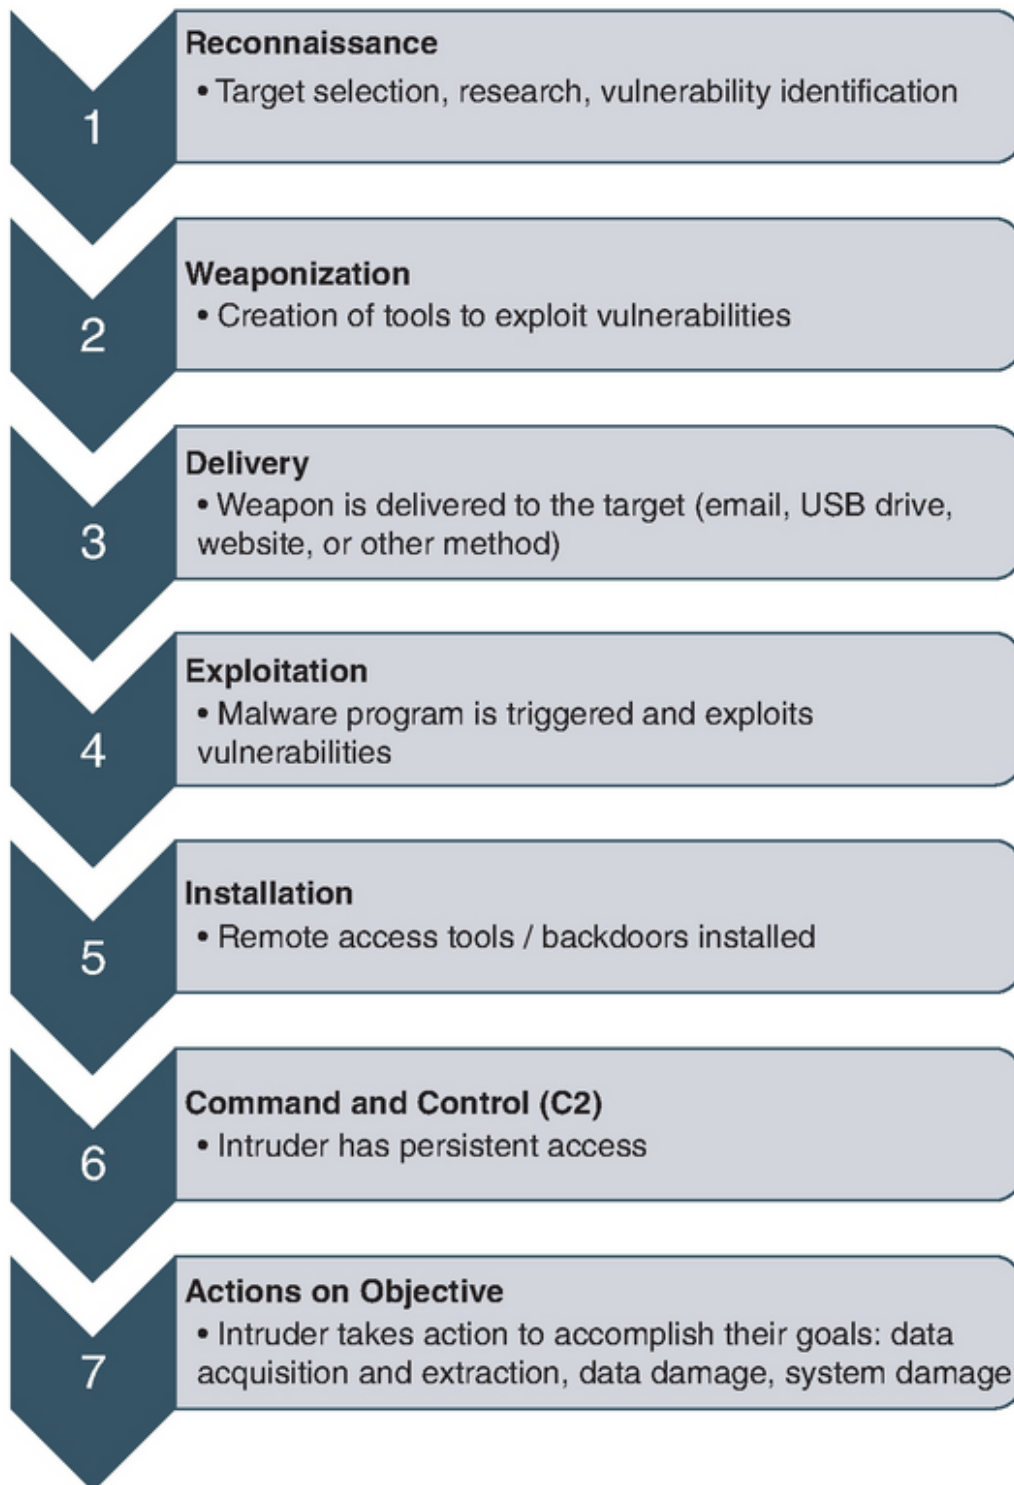


Figure 20.2 The Cyber Kill Chain.

The seven stages of the Cyber Kill Chain are as follows:

Reconnaissance involves target selection, research, and vulnerability identification. Adversaries gather intelligence about the target through both open-source research and direct scanning.

Weaponization involves creating tools to exploit the identified vulnerabilities. This may include building a malicious payload or combining malware with an exploit.

Delivery occurs when the weapon is transmitted to the target, such as through a phishing email, a USB drive, or a compromised website.

Exploitation uses a technical or human vulnerability to gain access. This may involve triggering a malware program or exploiting an unpatched software flaw.

Installation establishes persistent access by installing remote access tools or backdoors on the compromised system.

Command and Control (C2) establishes a communication channel that allows the attacker to remotely control the compromised system and issue commands.

Actions on Objective is the final stage, where the attacker accomplishes their mission goal, which may include data exfiltration, data destruction, or other malicious activities.

Diamond Model

The *Diamond Model of Intrusion Analysis* examines cyberattacks by analyzing the relationships between four core elements. In the Diamond Model, each observed malicious activity is treated as an event. Every event has four core features, represented as the vertices of a diamond:


The *adversary*, which represents the individual, group, or organization responsible for the malicious activity

The *capability*, which represents the tools, malware, exploits, or techniques the adversary used to carry out the attack

The *infrastructure*, which represents the systems, networks, domains, or services the adversary used to deliver, control, or support the attack

The *victim*, which represents the targeted organization, system, user, or data affected by the malicious activity

The Diamond Model is particularly useful for attribution: understanding who conducted an attack and how they operate. By examining the relationships among these four elements, analysts can use knowledge of one vertex to discover information about the others. For example, identifying the infrastructure used in an attack can lead investigators to the adversary behind it. [Figure 20.3](#) shows an example of a Diamond Model analysis of a compromised system.



[Figure 20.3](#) A Diamond Model analysis of a compromised system.

Exam Essentials

Logs provide crucial insights into security events by enabling accountability, traceability, and auditability. Security information and event management (SIEM) systems centralize log collection and use correlation to detect patterns of malicious activity across an organization.

Security event triage is the process of classifying, validating, and prioritizing security alerts. Analysts consider event severity, asset criticality, and potential business impact when determining which alerts require immediate investigation. Incident use cases help analysts recognize and respond to common attack patterns.

Threat actors vary in their intent, motivations, capabilities, and resources. Nation-state actors conduct espionage and sabotage with sophisticated resources. Organized cybercriminals seek financial gain. Hacktivists pursue political or social agendas. Insider threats exploit legitimate access. Script kiddies use publicly available tools with limited sophistication.

Cyber threat intelligence (CTI) involves collecting and analyzing information about threats, threat actors, and their tactics, techniques, and procedures (TTPs). Indicators of compromise (IoCs) are forensic artifacts that suggest a system has been breached. Organizations obtain CTI from open sources, commercial vendors, government agencies, and industry ISACs.

Threat frameworks provide structured approaches for analyzing attacks. The MITRE ATT&CK framework catalogs adversary tactics and techniques. The Cyber Kill Chain models attacks as seven sequential stages. The Diamond Model analyzes attacks through four elements: the adversary, the capability, the infrastructure, and the victim.

Practice Question 1

You've just been assigned to improve the existing log management system in your organization. Which solution would best aid in centralizing, protecting, and analyzing your log entries?

- A. DBMS
 - B. NIDS
 - C. SIEM
 - D. IAM
-

Practice Question 2

As a network security analyst, you are tasked with investigating an unusual surge in network traffic. You are looking for a way to connect the dots between different events and identify any potential security threats. Which attribute of effective logging is most useful for this task?

- A. Accountability
 - B. Auditability
 - C. Availability
 - D. Traceability
-

Practice Question 1 Explanation

A database management system (DBMS) is used primarily to manage databases and is not specifically designed for centralized log management. An identity and access management (IAM) system is more focused on managing user identities and controlling their access to resources.

While a network intrusion detection system (NIDS) does analyze network traffic and create logs of suspicious activity, it doesn't have the capability to centralize and manage logs from various sources.

The security information and event management (SIEM) system, however, does exactly what the question requires. It provides centralized logging capabilities, securely stores logs, and employs advanced analytics to scrutinize log data for potential security incidents. So, the most suitable solution to improve the log management system in this scenario would be a SIEM system.

Correct Answer: C. SIEM

Practice Question 2 Explanation

When you face an unusual event or a security incident, the ability to identify and track all related events is crucial. This process is known as traceability. It helps analysts connect the dots between various events and provides a comprehensive view of an incident.

Accountability, or identity attribution, while important, is more about determining the origin or cause of an event, which isn't the primary need in this scenario.

Auditability refers to documenting events, which is valuable but not the main requirement in this situation.

Availability is more closely related to the accessibility and uptime of systems and data, rather than to the investigation of security incidents.

Correct Answer: D. Traceability

CCSM *Certified in Cybersecurity Study Guide*, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 21 Incident Response: Objective 5.3 Understand Incident Response (IR)

While we strive to protect our systems and information against a wide variety of threats, the grim reality is that no matter how many controls we put in place, there's still a possibility that we'll fall victim to a security incident. Cybersecurity professionals must have plans and procedures to handle these incidents when they arise.

In this chapter, you'll learn about CC objective 5.3. The following subobjectives are covered in this chapter:

Data handling policy implementing Incident Response Plan (IRP)

Incident Response (IR) exercises (e.g., testing, tabletop)

Creating an Incident Response Program

This section explores the *incident response* process, focusing on a standard incident response process endorsed by the National Institute of Standards and Technology (NIST), as shown in [Figure 21.1](#). It uses the following phases:

Preparation, which includes the activities used to put together an incident response plan and team

Detection & Analysis, which identifies that an incident is taking place and determines the extent of the incident's impact

Containment, Eradication & Recovery, which limits the damage caused by an incident, removes the effects of the incident, and restores normal operations

Post-incident Activity, which analyzes the response process and identifies lessons learned to improve future response efforts




Figure 21.1 NIST incident response life cycle.

This process is fully described in the *NIST Computer Security Incident Handling Guide* (NIST SP 800-61, Revision 2). This document is widely used as a standard reference throughout the cybersecurity field.

Every organization should develop a cybersecurity incident response plan that outlines the policies, procedures, and guidelines it will follow when an incident occurs. This preparation process is extremely important because it provides structure and organization during a crisis.

The Importance of Planning

I've been involved in many security incidents throughout my career. In hindsight, it's clear that all the organizations that handled incidents well had one thing in common: They had thought through their incident response process and documented it in advance.

On the other hand, when I think about the incidents that didn't go very well, they typically occurred in organizations that didn't conduct prior planning. In those organizations, I commonly heard the sentiment, "Well, we're good at crisis management, and a security incident isn't that likely. We'll figure out the details when it happens."

That seat-of-the-pants approach to cybersecurity incident handling is a recipe for failure. The reality is that people make bad decisions in the heat of a crisis. Developing an incident response plan in advance of an incident allows you to make decisions in the calm environment of the planning phase, which then helps you exercise good judgment in the heat of an incident.

A formalized incident response plan should include several common elements:

It should begin with a *statement of purpose*. Why is the organization creating an incident response plan, and what is its scope? Which types of incidents does the plan cover? For example, is it restricted to only cybersecurity incidents, or does it cover any loss of sensitive information?

It should describe *clear strategies and goals* for the incident response effort. What are the highest priorities for first responders and those handling an incident at a more strategic level? If responders should prioritize containment over evidence preservation, make sure that's clear in the plan.

It should describe the *organization's approach to incident response*. Who is responsible for incident handling? What authority do they have?

It should also cover *communication* within the team, with other groups within the organization, and with third parties.

It should include the *senior management's approval*. You might need that authority when taking unpopular actions during incident response. If you can

point to the plan and show an irate administrator that the policy requiring disconnection of a system was signed by the CEO, that goes a long way.

Your incident response plan should also include a data-handling policy that outlines how the organization will manage and protect data throughout the incident response process. This policy should address how evidence is collected, preserved, and stored; how sensitive data discovered during an investigation is handled; and how data breach notification requirements will be met. Integrating data handling procedures into your IRP ensures that the organization meets its legal and regulatory obligations while maintaining the integrity of incident-related evidence.

As you develop your plan, you should consult NIST SP 800-61 to help guide your decisions. You also might find it helpful to look at some plans developed by other organizations. Of course, you won't be able to simply take someone else's plan and apply it to your organization, but it's always helpful to have a starting point. Many cybersecurity professionals have put countless hours into developing strong incident response plans, and there's no need to reinvent the wheel.

Building an Incident Response Team

One of the most important tasks that you'll undertake in your incident response program is to build and staff your incident response team.

This team will likely need to be available on a 24/7 basis, and you should have primary and backup personnel assigned to cover vacations as well as extended periods of operation. Incident handling is a wonderful professional development opportunity and helps team members keep their technical skills sharp.

Team Composition

Some of the groups that should be represented in your incident response team include the following:

- Management

- Information security personnel

Physical security team members

Technical subject matter experts, such as database administrators, developers, systems engineers, and virtualization experts

Legal counsel

Public relations and marketing staff

Human resources team members

Including the right team members is critical to building the relationships that you'll need during an incident. You won't necessarily need to activate all team members for any given incident, but each of these groups should have representatives trained and ready to participate before an incident strikes.

As you build out your incident response team, you may find that your organization lacks some of the capacity to handle security incidents. For example, you might discover that you don't have the forensic capabilities within your team to conduct investigations in support of incident response efforts. In those cases, you may want to consider retaining an external incident response provider to assist.

Exam Tip

You don't want to try to locate and negotiate a contract with a provider in the middle of an incident. Plan in advance and get the paperwork in place to use a provider immediately when you discover a problem. Your incident response team will be a crucial asset as you work to address the impact of a security incident. Be sure that you take the time now to design and train your team so that they're ready to respond in the event of an actual cybersecurity incident.

Training and Testing

Once your team is in place, you should work with them regularly. Don't wait until incidents occur to pull everyone together. Provide them with your incident response plan documentation and conduct regular training and testing

to ensure that they work well together and are ready to react quickly in the event of an incident.

Incident Response Exercises

Creating an incident response plan and building a team are important first steps, but organizations must also regularly *test* their incident response capabilities. Incident response exercises validate that plans work as expected, that team members understand their roles, and that communication channels function properly. Without regular testing, organizations may discover critical gaps in their plans only when a real incident occurs: the worst possible time to learn that something is broken.

The simplest form of IR exercise is the *tabletop exercise*. In a tabletop exercise, key stakeholders and incident response team members gather together and walk through a simulated incident scenario in a discussion-based format. A facilitator presents the scenario, which might involve a ransomware attack, a data breach, or a denial-of-service (DoS) incident, and then guides the group through each phase of the response.

Participants discuss what actions they would take, who they would notify, and how they would coordinate their efforts. Tabletop exercises are relatively inexpensive to conduct and don't require any changes to production systems, making them an accessible starting point for organizations of any size.

More advanced organizations may also conduct *simulation exercises*, sometimes called functional or operational exercises. In a simulation exercise, the incident response team actively responds to a realistic but controlled incident in real time. For example, a red team might launch a simulated attack against the organization's network while the incident response team detects, analyzes, and responds to the activity using their normal tools and procedures.

Simulation exercises provide a much more realistic test of the team's capabilities than tabletop exercises, but they also require more planning, resources, and coordination to execute safely without disrupting normal business operations.

Every IR exercise should conclude with a lessons-learned session. The team should document what went well, what problems arose, and what changes should be made to the incident response plan. These findings should then be

incorporated into an updated version of the plan, creating a cycle of continuous improvement.

Incident Communications Plan

One of the critical components of your incident response program is an *incident communications plan* that covers both internal and external communications. A good incident communications plan is a crucial element of stakeholder management.

Internal Communications

Incident notification and escalation procedures help ensure that the appropriate people within your organization know about an incident at the right time and are provided with the right information.

External Communications

Communicating with individuals and groups outside of your organization can be a much trickier task. Clearly, you want to make sure that you're limiting the communication of sensitive information to trusted parties. This is particularly important when there might be public or media interest in an incident. If word leaks out without approval, the incident might wind up in the news before your public relations team is ready to handle it. This might also jeopardize the integrity of your investigation by alerting attackers to the fact that you've discovered the incident and an incident response effort is underway.

In most cases, you aren't under a legal obligation to report security incidents to law enforcement, and the decision to do so is complex. Once you file a report with law enforcement, it's likely that details of the incident will become public, which may be undesirable. Also, law enforcement officials are held to much higher standards in gathering and processing evidence. Of course, you should always contact law enforcement if you think there is a threat to safety or you have a legal obligation to report a specific kind of incident.

Your legal team should be included in your incident response planning efforts. They should provide you with specific guidance about any laws or regulations

that apply to your organization and may require notification in the event of a security incident. For example, most states have privacy laws that require the timely notification of individuals when there is a compromise of personal information. You may also have obligations under other laws and regulations to notify government agencies, private regulatory bodies, customers, or the public about specific types of incidents, depending on their impact and the types of information involved.

Secure Communications

Your communications plan should not only describe *who* you will communicate with during an incident but should also describe *how* you will communicate. Make sure that you have secure communications paths in place, before an incident occurs, that provide you with confidential mechanisms to share information with trusted employees and third parties. Using secure channels prevents the inadvertent release of information to the public or adversaries.

Incident Identification and Response

Once you have an incident response plan in place and a team prepared, the incident response process enters a state of perpetual monitoring: watching for signs that an incident is taking place or has already occurred. There are many different ways that an organization might identify a security incident.

Security Data Sources

The key to successful incident identification is to have a robust security monitoring infrastructure. Data is crucial to incident detection and organizations have a responsibility to collect, analyze, and retain security information.

Many information sources can contribute data crucial to identifying and analyzing a possible security incident, including:

- Intrusion detection and prevention systems

- Firewalls

Authentication systems

System integrity monitors

Vulnerability scanners

System event logs

NetFlow connection records

Antimalware packages

If IT systems do one thing well, it's generating massive amounts of log information!

Correlating Security Information

Security professionals are responsible for collecting and correlating log information. Unassisted, that's almost an impossible undertaking. Fortunately, *security information and event management (SIEM)* technology can assist with this task. SIEM systems act as centralized log repository and analysis solutions. Security professionals can take the firehose of data they receive from security-related logs and point it at the SIEM, which can then do the heavy lifting of analysis. SIEM systems can detect possible incidents based on rules and algorithms, bringing them to the attention of security administrators for further review. They also provide a critical centralized information source to investigators pursuing a security incident.

Receiving Incident Reports

I've just discussed many of the ways that a security team might identify incidents based on internally generated data, and, in the best case, that's the way security professionals detect incidents—by noticing the signs of an incident as it occurs or shortly thereafter. Unfortunately, sometimes those monitoring systems fail to detect an incident, and we first learn of a security compromise by hearing from employees, customers, or external organizations who see the signs of a breach.

This might occur because a customer sees their personal information posted on the web, because a system on the corporate network begins attacking an external site due to commands received from a botnet, or because an employee

notices being unable to log in to an email account. The incident response team should have a consistent method for receiving, recording, and evaluating these reports.

Responding to Incidents

When a security professional identifies a potential incident, it's time to swing into incident response mode. The team member who first notices an incident and others who may be on duty have special first responder responsibilities. Just as in a medical emergency, the first person on the scene can have a tremendous impact on the successful response to an incident by acting quickly and decisively to protect the organization.

First responders should act quickly to contain the damage from a security incident. If they suspect that a system or group of systems may be compromised, they should isolate that system from the remainder of the network to contain the damage.

Depending on the technical circumstances, they may quarantine the system by removing it from the network, keeping it running to preserve evidence but cutting off the potentially compromised system's ability to communicate with attackers or infect other systems on the corporate network.

Exam Tip

This is a favorite topic for exam questions. Remember that a first responder's highest technical priority is to contain the damage by isolating affected systems.

As you build out your incident response capabilities, be sure to integrate them with your threat intelligence program. Organizations with strong strategic intelligence programs will be able to more rapidly and effectively identify potential incidents.

At the same time, remember that your adversaries are gathering intelligence on your organization and its operations. Counterintelligence programs are designed to thwart these efforts by denying adversaries access to intelligence or deliberately feeding them misinformation.

Exam Essentials

The four phases of the incident response process are preparation, detection & analysis, containment, eradication & recovery, and post-incident activity.

Incident response teams commonly include members from management, information security, physical security, legal counsel, public relations, and human resources, along with technical subject matter experts.

The highest technical priority of a first responder during incident response is to quickly contain the damage caused by the security incident.

Security information and event management (SIEM) systems act as centralized log repository and analysis solutions.

Practice Question 1

Jason is monitoring his organization's SIEM system watching for signs of unusual activity. Which phase of the NIST incident response process best describes his work?

- A. Preparation
- B. Detection & analysis
- C. Post-incident activity
- D. Containment, eradication & recovery

Practice Question 2

During his monitoring work, Jason identifies a high-priority security incident in progress. What should be his first priority?

- A. Identifying lessons learned
- B. Notifying senior management
- C. Recovering normal operations
- D. Containing the damage

Practice Question 1 Explanation

Jason is monitoring his environment for signs of unusual activity. This is an activity designed to identify security incidents and, therefore, belongs to the detection & analysis phase of the incident response process.

The preparation phase includes the activities used to put together an incident response program and team.

The containment, eradication, and recovery phase limits the damage caused by an incident, removes the effects of the incident, and restores normal operations.

The post-incident activity phase analyzes the response process and identifies lessons learned to improve future response efforts.

Correct Answer: B. Detection & analysis

Practice Question 2 Explanation

This is a tricky question because Jason and the incident response team will likely take all of these actions during the incident response effort. The key to answering this question correctly is to notice that it is asking for the first priority.

The first priority of any incident response effort should always be to contain the damage. Jason should immediately take steps to isolate any affected systems to stop the spread of the incident.

Notifying senior management and activating the incident response process is also a high priority, but Jason should first take immediate action to contain the damage. If other team members are available, he can ask them to perform notification and escalation while he is containing the damage.

Recovering normal operations and identifying lessons learned come much later in the process, after the initial emergency has been resolved.

Correct Answer: D. Containing the damage

CCSM *Certified in Cybersecurity Study Guide*, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 22 Asset Protection: Objective 5.4 Understand Asset Protection

Protecting an organization's technology assets is a fundamental responsibility of cybersecurity professionals. This includes managing assets throughout their lifecycle, maintaining secure configurations, and controlling changes to the technology environment.

In this chapter, you'll learn about CC objective 5.4. The following subobjectives are covered in this chapter:

Asset lifecycle management (e.g., End Of Life (EOL) software and devices)

Configuration and change management

Asset Lifecycle Management

Every technology asset has a lifecycle that begins when the organization acquires it and ends when the asset is retired and disposed of. Managing this lifecycle effectively is essential for maintaining security, because assets that are not properly tracked and maintained can become vulnerabilities that attackers exploit.

The asset management lifecycle consists of six key stages, as shown in [Figure 22.1](#).




Figure 22.1 The asset management lifecycle.

The six stages of the asset management lifecycle are:

Planning & Procurement. In this first stage, the organization identifies its need for a new asset and acquires it. Security considerations should be part of

the procurement process, ensuring that new assets meet the organization's security requirements before they are purchased.

Deployment & Configuration. Once an asset is acquired, it must be deployed to the environment and configured in accordance with the organization's security standards. This includes hardening the system, applying security baselines, installing required software, and registering the asset in the organization's asset inventory.

Operation & Maintenance. During the longest stage of the lifecycle, the asset is in active use. The organization must maintain the asset by applying patches and updates, monitoring its performance and security, and ensuring that it continues to meet the organization's needs.

Monitoring & Assessment. Throughout the operational phase, the organization should continuously monitor the asset for security issues and periodically assess whether it continues to meet security and business requirements.

End of Life (EOL) & End of Support (EOS). Eventually, every asset reaches the end of its useful life. This may happen when the manufacturer stops providing security updates (end of support) or when the asset can no longer meet the organization's needs.

Disposal & Decommissioning. In the final stage, the asset is removed from service and disposed of securely. This includes wiping or destroying any data stored on the asset, removing it from the asset inventory, and disposing of the physical hardware in an environmentally responsible manner.

Exam Tip

The CC exam specifically mentions end-of-life (EOL) software and devices as examples of asset lifecycle management. Be sure you understand that EOL assets no longer receive security updates from their manufacturer and pose a significant security risk if they remain in the environment.

Asset Inventory

Maintaining an accurate *asset inventory* is the foundation of effective asset lifecycle management. The asset inventory is a comprehensive record of all technology assets owned or managed by the organization, including hardware, software, network, cloud, and data assets. Without an accurate inventory, the organization cannot effectively protect its assets because it doesn't know what it has or where it is located.

The asset inventory should include key information about each asset, such as:

Asset identifier, type, and description

Serial number or license key

Physical location or logical assignment

Owner or responsible party

Vendor, date of acquisition, and expected EOL

Software version and patch level

Security classification of data stored on or processed by the asset

Organizations should regularly audit their asset inventory to ensure that it remains accurate and up to date.

End-of-Life Risks

One of the most significant security challenges in asset lifecycle management is dealing with *end-of-life (EOL)* and *end-of-support (EOS)* software and devices. When a manufacturer declares a product EOL, it means that the product is no longer being actively developed or sold. When a product reaches the EOS, the manufacturer will no longer provide security patches or technical support.

Running EOL or EOS software and devices in a production environment creates serious security risks:

Newly discovered vulnerabilities will not be patched, leaving the organization permanently exposed to those threats.

Compliance frameworks and regulations may prohibit the use of unsupported software, leading to audit findings or regulatory penalties.

Integration issues may arise when newer systems and applications are not compatible with outdated technology.

Insurance providers may not cover incidents that result from the use of unsupported technology.

Organizations should maintain a list of EOL and EOS dates for all critical software and hardware in their environment and develop migration plans well in advance of those dates. When it is not possible to immediately replace an EOL asset, the organization should implement compensating controls, such as network segmentation, enhanced monitoring, and additional access restrictions, to mitigate the risk.

Configuration Management

Configuration management is the process of establishing and maintaining the settings and configurations of an organization's technology assets. It ensures that systems are set up correctly, remain in a known and secure state, and that any changes to their configuration are tracked and controlled.

Baselines

Baselining is an important component of configuration management. A baseline is a snapshot of a system or application at a given point in time. It may be used to assess whether a system has changed outside of an approved change management process. System administrators can compare a running system to a baseline to identify all changes to the system and then compare those changes to approved changes.

Security baselines define the minimum security settings that should be applied to a particular type of system or device. When a new server is deployed, it is configured according to this baseline to ensure that it meets the organization's security requirements from the start.

Patch and Update Management

Patch management ensures that systems and applications receive all security updates issued by manufacturers to address known vulnerabilities. An

effective patch management program includes several key activities:

Monitoring vendor announcements and security advisories for new patches

Evaluating the relevance and urgency of each patch based on the organization's environment

Testing patches in a non-production environment before deploying them to production systems

Deploying approved patches in a timely manner, prioritizing critical security patches

Verifying that patches have been successfully applied and have not introduced new issues

Configuration Vulnerabilities

Configuration vulnerabilities can have serious impacts on enterprise security. Common configuration vulnerabilities include:

Default configurations. Taking a system directly from a manufacturer and installing it on the network without modifying the default configuration can leave open ports, default passwords, guest accounts, and other security issues in place.

Weak security settings. Systems that are misconfigured or configured with weak security settings can allow attackers to gain control of the device.

Cryptographic weaknesses. If an administrator inadvertently configures weak cipher suites or weak protocol implementations, communications to and from that device might be subject to eavesdropping and tampering.

Change Management

Change management is the process of controlling modifications to an organization's technology environment. It ensures that changes are made in a planned, controlled manner that minimizes the risk of disruption and maintains the security of the environment.

The Change Management Process

A well-designed change management process typically includes the six steps shown in [Figure 22.2](#):

Request. The process begins when someone submits a request for a change, describing what change is being proposed, why it is needed, and what systems or users will be affected.

Evaluation and Impact Assessment. The change is evaluated to determine its potential environmental impact, including assessing risks and identifying dependencies.

Approval. The change is reviewed and approved (or denied) by the appropriate authority. For significant changes, the change may need to be reviewed by a change advisory board (CAB).

Implementation. Once approved, the change is implemented during a planned maintenance window, when possible, as documented.

Documentation. The change and its results are documented for future reference and become part of the organization's change history.

Rollback. If the change causes unexpected problems, a *rollback plan* restores the previous configuration. Developing and testing rollback plans before implementing a change is essential.

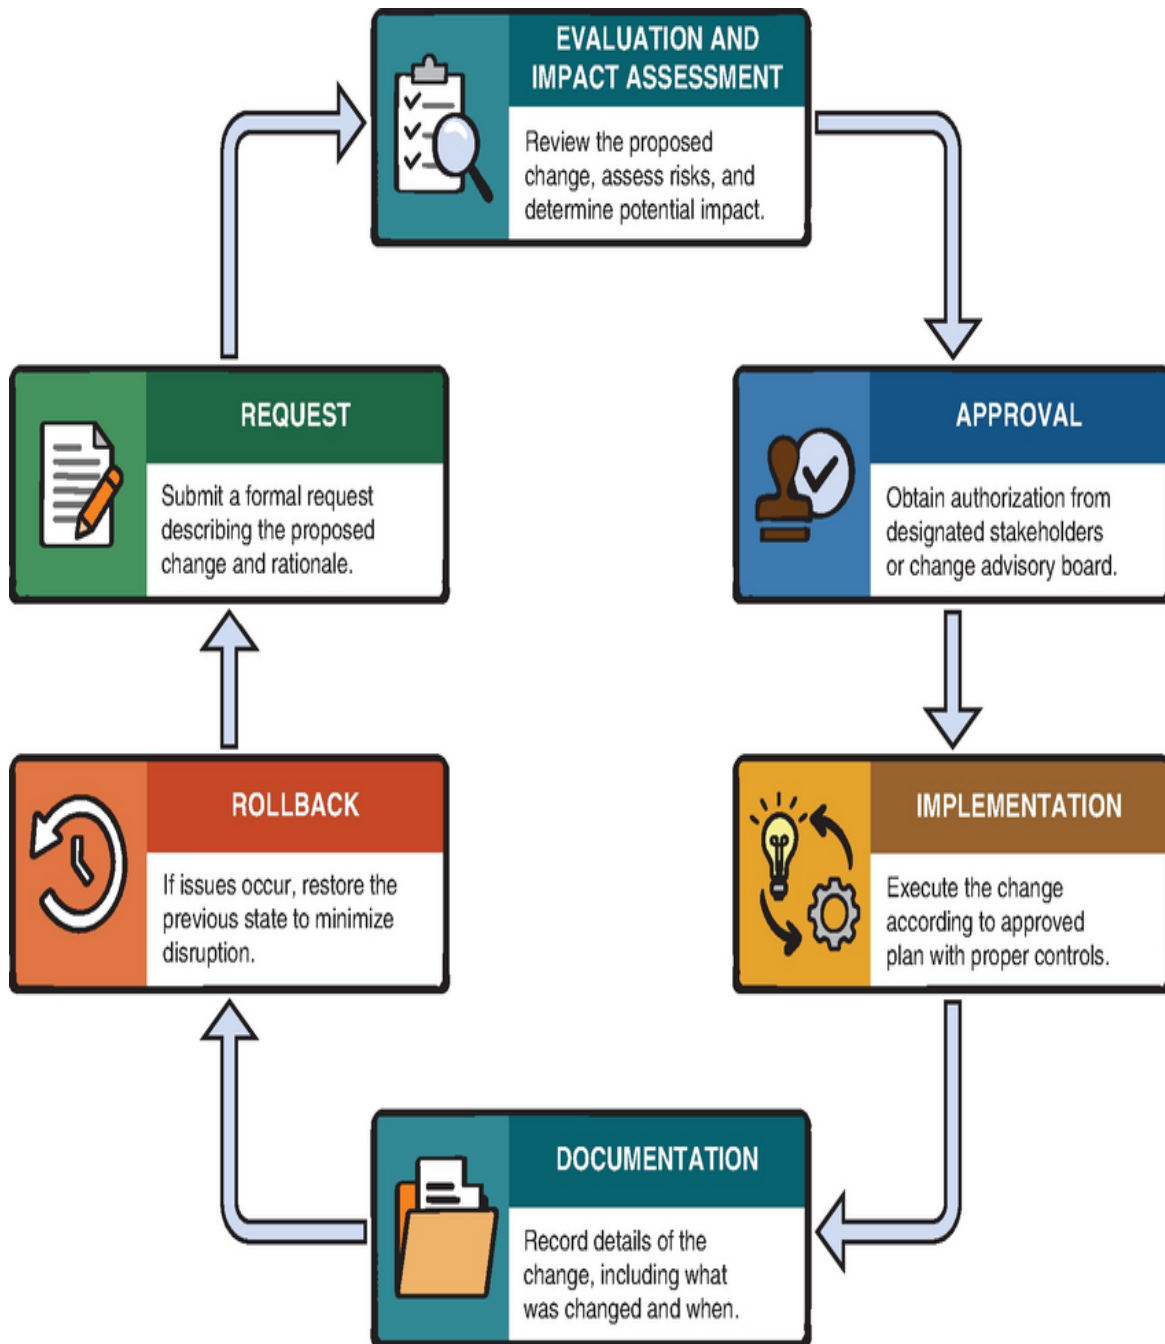


Figure 22.2 The change management process.

Emergency Changes

Not all changes can wait for the full change management process. When a critical security vulnerability is discovered and an emergency patch must be applied immediately, the organization needs a process for handling emergency

changes. Emergency change procedures typically allow for expedited approval and implementation but still require documentation after the fact.

Exam Essentials

Assets progress through stages, including planning and procurement, deployment and configuration, operation and maintenance, monitoring and assessment, end of life, and disposal.

End-of-life (EOL) and end-of-support (EOS) software and devices no longer receive security patches. Running unsupported technology creates serious security risks and may violate compliance requirements.

Configuration management establishes and maintains the settings of technology assets. Key components include baselines and patch management.

Change management ensures that modifications are properly requested, evaluated, approved, implemented, documented, and reversible through rollback plans.

Practice Question 1

What is the primary risk of running end-of-life (EOL) software in a production environment?

- A. The software will run more slowly than newer versions.
- B. Newly discovered vulnerabilities will not be patched by the manufacturer.
- C. The software will automatically stop functioning on the EOL date.
- D. Users will not be able to save files created with the software.

Practice Question 2

Your organization is implementing a change management policy. Which of the following is an essential element that should be developed before implementing any change?

- A. A marketing plan for the change
 - B. A rollback plan to restore the previous configuration
 - C. A new budget for the IT department
 - D. A press release announcing the change
-

Practice Question 1 Explanation

The primary risk of running EOL software is that newly discovered vulnerabilities will not be patched by the manufacturer, leaving the organization permanently exposed to those threats.

EOL software does not automatically stop working, nor does it necessarily run slower. There should be no impact on the ability of users to save files.

Correct Answer: B. Newly discovered vulnerabilities will not be patched by the manufacturer.

Practice Question 2 Explanation

A rollback plan is essential for change management because it enables restoring the previous configuration if the change causes unexpected problems.

Marketing plans, budgets, and press releases are not standard elements of the change management process.

Correct Answer: B. A rollback plan to restore the previous configuration

CCSM Certified in Cybersecurity Study Guide, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 23 Security Testing: Objective 5.5

Understand Security Testing

Security testing is the process of evaluating an organization's security controls and defenses to identify vulnerabilities before attackers can exploit them.

Regular security testing is an essential part of a mature cybersecurity program.

In this chapter, you'll learn about CC objective 5.5. The following subobjectives are covered in this chapter:

Security readiness testing (e.g., blue teaming, purple teaming, red teaming)

Application testing (e.g., vulnerability scanning, static analysis, dynamic analysis, threat modeling)

Physical penetration testing (e.g., phishing, tailgating, impersonation)

Security Readiness Testing

Security readiness testing evaluates an organization's ability to detect, prevent, and respond to incidents. These tests go beyond simple vulnerability scanning by simulating realistic attack scenarios and measuring the organization's ability to defend against them.

Security readiness testing typically involves specialized teams that play different roles in the testing process. The three most common team-based approaches are red teaming, blue teaming, and purple teaming.

Red Teaming

A *red team* is a group of security professionals who simulate the actions of a real-world attacker. The goal of red teaming is to test the organization's defenses by attempting to penetrate its security controls using the same tactics, techniques, and procedures (TTPs) that actual attackers would employ.

Red team exercises are typically conducted without the knowledge of the organization's security staff, thereby testing the organization's ability to detect and respond to an actual attack. Red team members may use a variety of techniques, including:

Attempting to gain unauthorized access to systems and networks through technical exploits

Using social engineering techniques to trick employees into revealing credentials or granting access

Testing physical security by attempting to gain unauthorized access to facilities

Simulating advanced persistent threats (APTs) that target specific high-value assets

At the conclusion of a red team exercise, the team produces a detailed report documenting the vulnerabilities they identified, the techniques they used, and recommendations to improve the organization's defenses.

Blue Teaming

The *blue team* represents the organization's defensive security personnel. In contrast to the red team, which plays the role of the attacker, the blue team is responsible for detecting, preventing, and responding to attacks. In many organizations, the blue team comprises security operations center (SOC) analysts, incident responders, and other security staff who defend the environment on a daily basis.

During a security readiness exercise, the blue team works to detect the red team's activities, contain any simulated breaches, and respond to the attack using the organization's existing tools and procedures.

Purple Teaming

Purple teaming is a collaborative approach that combines elements of both red and blue teaming. Rather than operating independently, the red and blue teams work together in a purple-team exercise to maximize learning opportunities for both sides.

In a purple-team exercise, the red team demonstrates specific attack techniques while the blue team observes and practices detecting and responding to those techniques. After each attack scenario, both teams discuss what happened, identify gaps in detection and response, and work together to develop improved defenses.

Exam Tip

Remember the three team colors: red teams attack, blue teams defend, and purple teams combine both approaches in a collaborative exercise. The CC exam may ask you to distinguish between these different testing approaches. If you get confused, remember the “color trick”: red + blue = purple!

Application Testing

Application testing focuses on identifying security vulnerabilities in software applications and systems. This is a critical area of security testing because applications are often the primary target for attackers seeking to gain access to sensitive data or compromise an organization’s systems.

Vulnerability Scanning

Vulnerability scanning is the automated process of probing systems, networks, and applications to identify known security weaknesses. Vulnerability scanners use databases of known vulnerabilities to check whether systems are susceptible to specific attacks.

Vulnerability scans can be classified based on the scanner’s perspective:

Authenticated scans use valid credentials to log into systems and perform a more thorough assessment from the inside. These scans can identify vulnerabilities that would not be visible to an external attacker.

Unauthenticated scans test systems from the outside without credentials, simulating the perspective of an external attacker who has not yet gained

access.

Organizations should conduct vulnerability scans regularly and after any significant changes to their environment.

Static Analysis

Static analysis is a technique for examining an application's source code or compiled binary without actually executing the program. Static analysis tools scan the code for patterns that indicate potential security vulnerabilities, such as buffer overflows, injection flaws, and improper error handling.

The primary advantage of static analysis is that it can be performed early in the software development life cycle, before the application is deployed. Many organizations integrate static analysis tools into their continuous integration and continuous delivery (CI/CD) pipelines to automatically scan code for vulnerabilities as it is written.

Dynamic Analysis

Dynamic analysis tests applications while they are running. Unlike static analysis, which examines the code itself, dynamic analysis interacts with the application as a user would, sending inputs and observing the application's behavior to identify vulnerabilities.

Dynamic analysis tools can identify vulnerabilities that static analysis cannot detect, such as:

Authentication and session management issues

Cross-site scripting (XSS) vulnerabilities

SQL injection flaws

Server configuration errors

Runtime errors that could be exploited by an attacker

Threat Modeling

Threat modeling is a structured approach to identifying and evaluating potential threats to an application or system. A typical threat modeling process involves several steps:

- . Identify assets that the application needs to protect.
- . Create an architecture overview documenting the application's components, data flows, trust boundaries, and external dependencies.
- . Identify threats using frameworks, such as MITRE ATT&CK (discussed in [Chapter 20, “Security Operations”](#)).
- . Evaluate and prioritize threats based on likelihood and potential impact.
- . Develop mitigations to address the identified threats.

Threat modeling is most effective when performed early in the development process, ideally during the design phase.

Exam Tip

Know the difference between static and dynamic analysis. Static analysis examines code without running it and can be done early in development. Dynamic analysis tests the running application by interacting with it. Both have strengths and limitations, and organizations often use both approaches together for comprehensive testing.

Physical Penetration Testing

Physical penetration testing evaluates the effectiveness of an organization's physical security controls by attempting to bypass them using techniques similar to those a real attacker might employ. The CC exam objectives specifically mention three common techniques: phishing, tailgating, and impersonation.

Phishing

Phishing is a social engineering technique in which an attacker sends deceptive messages, typically emails, designed to trick recipients into revealing sensitive information or performing actions that compromise security. In the context of physical penetration testing, phishing tests evaluate how well employees can recognize and resist these deceptive messages.

There are several variations of phishing that testers may use:

Spear phishing targets specific individuals or groups with personalized messages that appear to come from a trusted source.

Whaling targets senior executives or other high-value individuals within an organization.

Vishing (voice phishing) uses phone calls rather than emails to trick victims into revealing information.

Smishing uses text messages (SMS) to deliver phishing attacks.

Tailgating

Tailgating occurs when an unauthorized person follows an authorized person through a secured entrance without using their own credentials. Effective countermeasures against tailgating include security vestibules (mantraps) that only allow one person to pass at a time, security awareness training, and security guards who verify credentials at building entrances.

Impersonation

Impersonation involves an attacker pretending to be someone they are not in order to gain access to a facility or information. During a physical penetration test, the tester might impersonate a delivery person, a repair technician, an IT support specialist, or even a fellow employee to gain entry to a restricted area.

Effective defenses against impersonation include visitor management procedures, badge systems that clearly identify employees and visitors, and a security culture that encourages employees to verify the identity of unfamiliar individuals before granting access.

Exam Tip

Physical penetration testing uses the same social engineering principles as digital attacks but applies them in the physical world. Phishing targets employees through deceptive messages, tailgating exploits physical access controls, and impersonation relies on pretending to be an authorized person. All three techniques test the human element of security.

Exam Essentials

Red teams simulate attacks. Blue teams defend. Purple teams combine both in a collaborative exercise to improve detection and response capabilities.

Vulnerability scanning identifies known vulnerabilities. Static analysis examines source code without running it. Dynamic analysis tests the running application. Threat modeling identifies potential threats at the design level.

Vulnerability scanners probe systems for known weaknesses. Authenticated scans provide more thorough results than unauthenticated scans.

Phishing uses deceptive messages. Tailgating involves following an authorized person through a secured entrance. Impersonation involves pretending to be an authorized person.

Practice Question 1

Which team in a security readiness exercise simulates the actions of a real-world attacker?

- A. Blue team
 - B. Purple team
 - C. Red team
 - D. Green team
-

Practice Question 2

An unauthorized person follows an employee through a badge-controlled door without scanning their own badge. What type of attack is this?

- A. Phishing
 - B. Impersonation
 - C. Tailgating
 - D. Vishing
-

Practice Question 1 Explanation

The red team simulates the actions of a real-world attacker to test the organization's defenses. The blue team defends. The purple team combines both approaches. There is no standard "green team" in security readiness testing.

Correct Answer: C. Red team

Practice Question 2 Explanation

Tailgating (also called piggybacking) occurs when an unauthorized person follows an authorized person through a secured entrance without using their own credentials. Phishing uses deceptive messages, impersonation involves pretending to be someone else, and vishing is voice phishing.

Correct Answer: C. Tailgating

CCSM *Certified in Cybersecurity Study Guide*, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Chapter 24 Artificial Intelligence

Artificial intelligence (AI) is transforming the cybersecurity landscape at an extraordinary pace. Threat actors are leveraging AI to discover vulnerabilities, accelerate reconnaissance, and launch highly personalized social engineering attacks.

At the same time, AI is helping cybersecurity professionals automate repetitive tasks, improve threat detection, increase efficiency, and reduce human error. Security professionals must understand how AI intersects with the core principles of their discipline.

In this chapter, you'll learn about the AI security concepts that ISC2 has integrated across the CC exam domains. The following topics are covered in this chapter:

What is artificial intelligence?

AI security principles

AI and organizational resilience

AI and access controls

AI and network security

AI and security operations

What Is Artificial Intelligence?

Before diving into the security implications of AI, it's important to understand what the technology is and how it works. *Artificial intelligence (AI)* is a broad term for computer systems that can perform tasks that normally require human intelligence, such as recognizing patterns, making decisions, and generating content. AI systems range from simple rule-based automation to sophisticated models that learn from data.

Machine learning (ML) is a subset of AI in which systems learn from data rather than following explicitly programmed rules. An ML system is trained on a dataset and uses that experience to make predictions or decisions about

new data it has not seen before. Most of the AI applications you'll encounter in cybersecurity rely on machine learning.

A machine learning system is built around a *model*, the mathematical structure it uses to make predictions. During training, a model analyzes large amounts of data and gradually adjusts its internal parameters to produce more accurate results. The data used during this process is called the *training data*. Once training is complete, the model can be used in production to analyze new inputs, a process known as *inference*.

A *large language model (LLM)* is a type of AI system trained on massive amounts of text data. LLMs can generate human-like text, answer questions, summarize documents, and write code. LLMs are commonly available through chat interfaces, such as OpenAI's ChatGPT, Anthropic's Claude, and Google's Gemini. They are also available programmatically through application programming interfaces (APIs) that allow automated access to LLM technology.

AI is a double-edged sword in cybersecurity. On the defensive side, AI powers tools that detect intrusions, correlate security events, and identify anomalous user behavior. On the offensive side, threat actors are leveraging AI to craft convincing phishing messages, discover vulnerabilities faster, and automate reconnaissance. Understanding both sides of this equation is essential for any security professional.

AI on the CC Exam

In April 2026, ISC2 published the *Exam Guidance for Artificial Intelligence*, a document that maps how AI security concepts are incorporated into the exam outlines for all ISC2 certifications, including the CC. The guidance confirms that ISC2 has integrated foundational AI concepts across all five CC exam domains.

You do not need to be an AI expert to pass the CC exam. ISC2's approach is to test whether entry-level practitioners can identify AI assets, recognize automated threats, and support the governance frameworks that keep AI systems secure.

The topics in this chapter reflect the AI-related knowledge that ISC2 expects CC candidates to demonstrate. Think of AI as a lens applied across the domains you've already studied: the same principles of confidentiality, integrity, availability, least privilege, and defense in depth apply, but with new considerations for intelligent systems.

You can download the full ISC2 Exam Guidance for Artificial Intelligence from the ISC2 website at <https://www.isc2.org/Insights/2026/04/ISC2-Publishes-Exam-Guidance-AI>.

For now, ISC2's coverage of artificial intelligence topics is a bit disjointed. I expect that these topics will be fully integrated into the next version of the exam objectives, which will likely be released in 2029.

AI Security Principles

AI introduces new considerations for each of the three pillars of information security: confidentiality, integrity, and availability. Organizations that deploy AI systems must think carefully about how those systems handle sensitive data, whether their outputs can be trusted, and how to keep them running reliably. The same principles you learned earlier in this book apply to AI, but there are some important nuances.

Hallucination

Hallucination occurs when a large language model generates output that sounds plausible but is factually incorrect. LLMs produce text by predicting the most likely next word in a sequence, not by retrieving verified facts. This means they can fabricate information, cite nonexistent sources, or present outdated data with high confidence.

Hallucination poses a real security risk for organizations that rely on AI-generated content. An employee who trusts a hallucinated answer could make a flawed decision, misconfigure a security control, or share inaccurate information with customers, company leadership, or regulators.

Organizations should train staff to treat AI output as a starting point that requires human verification, not as a trusted source of truth. Review processes for AI-generated content are an important part of maintaining data integrity and reducing errors that could compromise the organization's security posture.

Data Integrity and Model Poisoning

One of the most significant risks to AI systems is *model poisoning*. In a model poisoning attack, an adversary tampers with the data used to train a machine learning model. Because the model learns its behavior from its training data, corrupted training data can lead to dangerously inaccurate results. For example, a poisoned spam filter might learn to let phishing emails through, or a poisoned malware detector might ignore a specific threat family.

Protecting data integrity throughout the AI system lifecycle is essential. Organizations should apply strict controls over who can access and modify training datasets, just as they would for any other critical information asset. Version control, file integrity monitoring, checksums, and access logging all help detect unauthorized changes before they compromise a model's output.

Ethical Use of AI

The ethical use of AI is an increasingly important component of a comprehensive security culture. Organizations should prioritize *transparency* in how their AI systems make decisions and guard against *algorithmic bias*,

which occurs when an AI system produces unfair or discriminatory results due to flaws in its training data or design.

Security professionals play a role in ensuring that AI adoption aligns with the organization's risk appetite and ethical standards. AI tools are subject to the same organizational policies and legal requirements as traditional software. Organizations should adopt AI-specific policies that address accountability, fairness, and the responsible handling of the data that powers these systems.

AI and Governance, Risk, and Compliance

AI systems must be incorporated into an organization's broader *governance, risk, and compliance (GRC)* program. This means treating AI tools as assets that require the same level of oversight as any other technology. Organizations should ensure that AI deployments comply with applicable laws and regulations and that senior leadership maintains visibility into how AI is being used across the enterprise.

Risk assessments should account for AI-specific threats, including model poisoning, prompt injection, data leakage, denial of service (DoS), and the potential for unexpected behavior. By integrating AI into existing GRC frameworks, organizations can ensure that they adopt AI responsibly and in a way that supports their overall security posture.

AI and Organizational Resilience

AI systems create new challenges for business continuity, disaster recovery, and incident response planning. Organizations that depend on AI-driven tools for critical operations must plan for the possibility that those tools will fail, degrade, or be attacked.

AI-driven Incident Detection

From a response perspective, AI-driven tools can help detect security incidents early. Machine learning models can analyze large volumes of log data and network traffic far more quickly than human analysts, flagging suspicious patterns that might otherwise go unnoticed.

Entry-level security professionals should understand how to follow established playbooks that now account for automated threats and how to provide support during the initial triage of a suspected breach.

Backing Up AI Systems

Recovery and continuity planning must extend beyond traditional data to include the specific configurations and datasets that power AI services. A machine learning model is only as good as its training data, and if that data is lost or corrupted, the model may need to be retrained from scratch. Organizations should ensure that their backup procedures cover model files, training datasets, and configuration parameters.

Model Drift

Model drift is a gradual decline in an AI system's performance over time. This can happen because the data the model encounters in the real world diverges from the data it was trained on. For example, a fraud detection model trained on last year's transaction patterns may become less effective as spending habits change.

Organizations should monitor their AI systems for signs of model drift and establish thresholds that trigger a review or retraining effort. If left unaddressed, model drift can silently degrade the quality of critical business decisions.

AI and Access Controls

Access control is the first line of defense for any system, and AI is no exception. However, AI introduces a new class of nonhuman entities that must be managed alongside traditional user accounts.

Managing AI Service Accounts

Just as human users are managed through a formal lifecycle, from provisioning to deprovisioning, AI bots and automated service accounts must

be managed through a formal lifecycle.

Organizations should apply the *principle of least privilege* to these accounts, ensuring that automated systems only have the permissions necessary to perform their designated tasks. When an AI tool is retired or replaced, its accounts should be promptly disabled to prevent unauthorized access.

Organizations should also conduct periodic reviews of these AI service accounts, just as they do for human user accounts. Regular access reviews verify that each account still serves a legitimate purpose, that its permissions remain aligned with its current function, and that accounts associated with retired or repurposed AI tools are promptly deprovisioned. Without scheduled reviews, AI service accounts can quietly accumulate excess privileges or persist long after they are needed, creating an attractive target for attackers.

AI-enhanced Authentication

AI is also being used to strengthen authentication mechanisms. Behavioral analysis powered by machine learning can detect anomalies in user login patterns, such as *impossible travel*, where a user appears to log in from two distant locations within an impossibly short time frame. For example, if a user logs in from New York and then appears to log in again from Tokyo 15 minutes later, the system can flag that event as suspicious and require additional verification before granting access.

These capabilities enhance existing multifactor authentication (MFA) systems by adding an intelligent layer of analysis. Rather than relying solely on static rules, AI-driven authentication can adapt in real time, adjusting the verification level based on the risk profile of each login attempt.

A user signing in from a familiar device at a normal time might pass through quickly, whereas an unusual login may trigger additional challenges. Security professionals should understand both how to secure AI's access to organizational systems and how to use AI to protect user identities across the enterprise.

AI and Network Security

AI is changing how organizations monitor and defend their networks. Machine learning models can analyze traffic patterns in ways that go far beyond traditional signature-based tools, identifying subtle anomalies that might indicate an attack in progress.

AI-powered Firewalls and Intrusion Detection

AI-powered firewalls and intrusion detection systems (IDS) use machine learning to identify unusual network behavior. Unlike traditional systems that rely on known signatures, these tools can detect novel threats by recognizing patterns that deviate from a learned baseline.

Entry-level practitioners should understand how these tools work and be prepared to monitor dashboards and report potential anomalies to senior analysts.

Network Segmentation for AI Environments

Organizations that develop or deploy AI systems should use *network segmentation* to isolate those environments. AI development environments should be isolated from sensitive production data to limit the damage if one environment is compromised. This foundational concept supports the implementation of Zero Trust (ZT) principles, which assume that no user, device, or service should be trusted by default, regardless of their location on the network.

By segmenting sensitive and/or risky AI workloads from the rest of the network, organizations can protect the high-value training data that powers their models and reduce the risk of lateral movement by attackers.

AI and Security Operations

Security operations teams are increasingly working alongside AI tools in their day-to-day activities. Understanding how these tools work and where they introduce new risks is essential for any security professional.

SIEM and AI-driven Correlation

Security information and event management (SIEM) systems collect and analyze security data across the enterprise. SIEMs are now using AI to correlate data from across the enterprise and reduce *alert fatigue*, the overwhelming volume of security notifications that can cause analysts to miss genuine threats. Machine learning models help distinguish between routine events and high-priority incidents that require human intervention.

Entry-level practitioners should understand how to handle the outputs of these automated systems and know when to escalate an alert for further investigation.

Securing the AI Workspace

The widespread availability of web-based AI tools and LLMs creates a new risk: *data leakage*. When employees interact with public AI services, they may inadvertently share sensitive organizational information with third-party providers. Source code, customer data, internal documents, and proprietary business strategies could all be exposed through casual use of these tools.

Once that information is submitted to an external AI service, the organization loses control over how it is stored, used, or retained. In some cases, submitted data may even be incorporated into the model's future training, making it potentially accessible to other users of the service.

Organizations should establish clear policies about the acceptable use of AI tools in the workplace. Security professionals serve as effective human firewalls, protecting the organization's data integrity in an increasingly automated world. Security awareness training should include guidance on the risks of sharing sensitive information with public AI services.

Exam Essentials

AI introduces new security considerations across all five CC domains. Model poisoning threatens data integrity, ethical AI use requires transparency and fairness, and AI systems must be included in governance, risk, and compliance (GRC) programs. Organizations should treat AI tools with the same level of oversight as any other technology asset.

Business continuity planning must account for AI systems. This includes backing up model files and training data, planning for AI-driven incident detection, and monitoring for model drift, where an AI system's performance gradually degrades over time.

AI entities require access controls just like human users. Automated service accounts should follow the principle of least privilege and be managed through a formal lifecycle. AI also enhances authentication through behavioral analysis and anomaly detection.

Network security for AI includes segmentation, Zero Trust (ZT), and AI-powered monitoring. AI development environments should be isolated from production data, and AI-powered firewalls and intrusion detection systems can identify threats that traditional tools miss.

Security operations teams work alongside AI tools daily. SIEM platforms use AI to reduce alert fatigue, and organizations must establish policies to prevent data leakage when employees use public AI services and large language models.

Practice Question 1

Ravi is concerned that an adversary may have tampered with the data used to train his organization's fraud detection model. Which type of AI-specific attack is Ravi most concerned about?

- A. Model drift
 - B. Data leakage
 - C. Model poisoning
 - D. Impossible travel
-

Practice Question 2

Keiko's organization recently deployed a machine learning model to assist with network threat detection. She wants to ensure that the model's service account follows security best practices. Which principle should she apply to the model's account permissions?

- A. Defense in depth
 - B. Separation of duties
 - C. Least privilege
 - D. Due diligence
-

Practice Question 1 Explanation

Ravi is concerned about model poisoning. In a model poisoning attack, an adversary tampers with the training data used by a machine learning model, causing it to produce inaccurate or dangerous results.

Model drift is a gradual decline in model performance over time as real-world conditions change. It is not caused by an adversary tampering with training data.

Data leakage occurs when sensitive information is inadvertently shared with unauthorized parties, such as through the use of public AI tools. This is not related to the integrity of training data.

Impossible travel is an authentication anomaly where a user appears to log in from two distant locations in an unrealistically short time. It is an access control concept, not an attack on AI training data.

Correct Answer: C. Model poisoning

Practice Question 2 Explanation

Keiko should apply the principle of least privilege to the model's service account. This means granting the account only the minimum permissions necessary to perform its designated tasks. Just as human users should, AI bots and automated service accounts should operate with the fewest possible privileges.

Defense in depth involves layering multiple security controls to protect against threats. While important, it does not specifically address account permissions for an AI service account.

Separation of duties ensures that no single individual can complete a critical task alone. This is an organizational control that does not directly apply to configuring permissions for an AI service account.

Due diligence involves conducting thorough research and investigation before making decisions. While relevant to AI deployment broadly, this principle does not govern account permissions.

Correct Answer: C. Least privilege

CCSM *Certified in Cybersecurity Study Guide*, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Index

Page numbers followed by *f* and *t* refer to figures and tables, respectively.

A

AAA (Authentication, Authorization, Accounting), [13](#)–23. *See also individual terms*

Access control:

in AI, [236](#)–237, [241](#)

biometrics, [20f](#)

in cloud storage, [165](#)–166

as data control, [156](#)

in data life cycle, [186](#)

digital, [14](#), [15f](#)

discretionary, [133](#), [136](#)

firewalls for, [158](#)

logical, *see* [Logical access control](#)

mandatory, [132](#)–133

microsegmentation for, [160](#)

models of, [132](#)–134

NFC technology for, [140](#)

physical, [15f](#), [229](#)

process of, [13](#)–16, [156](#)

role-based, [134](#)

in SaaS, [168](#)

Windows list, [133f](#)

Access logging, [63](#), [235](#). *See also* [Logging](#)

Access reviews, [57](#), [123](#), [125](#), [126](#), [129](#), [237](#)

Accountability, [27–28](#), [62](#), [66](#), [74](#), [193](#), [203](#), [235](#)

Accounting, [14](#), [15](#), [38](#), [132](#)

Actions on objectives, [199](#)

Active scanning, [197](#), [197f](#)

AD (Microsoft Active Directory), [128](#)

Administrative controls, [57](#), [59](#), [156](#)

AES (Advanced Encryption Standard), [51](#), [149](#), [177](#)

Agreement, notice, and communication, [28](#), [35](#)

AI (Artificial Intelligence), [233–241](#)

and access controls, [236–237](#), [241](#)

definition of, [233–234](#)

log monitoring by, [194](#)

and network security, [237–238](#)

organizational resilience of, [236](#)

and security operations, [238–239](#)

security principles of, [234–236](#)

AICPA (American Institute of Certified Public Accountants), [27](#)

Alert fatigue, [238](#)

Algorithmic bias, [235](#)

Anomaly detection, [239](#)

Anthropic, [234](#)

Antimalware software, [156](#), [211](#). *See also* [Malware](#)

Antivirus software, [56–57](#)

APIPA (Automatic Private IP Addressing), [144](#)
APIs (Application Programming Interfaces), [234](#)
Application controls, [156](#)
Application layer, [142](#)
Application testing, [227](#)–228
The Art of Intrusion (Mitnick), [105](#)
Asset inventory, [217](#)
Asset lifecycle management, [215](#)–218, [216f](#)
Asset protection, [215](#)–221
Asymmetric encryption, [177](#)
ATT&CK. *See also* [MITRE ATT&CK](#)
Audit management, [78](#)
Auditability, [193](#), [203](#)
Authenticated scans, [227](#)
Authentication:
as access control, [14](#)–15, [156](#)
AI-enhanced, [237](#)
biometric, [8](#), [18](#), [19](#), [20f](#)
captive portals for, [148f](#)
dynamic analysis of, [227](#)
enterprise, [148](#)
factors of, [18](#)–21
and impossible travel, [241](#)
information from, [211](#)
multi-factor, [20](#)–21, [38](#), [107](#), [156](#), [230](#), [237](#)
for privacy, [31](#)

- in provisioning, [125](#)
- of Wi-Fi users, [146](#)
- in WPA3, [149](#)
- Authority and trust, [105](#)
- Authorization, [131](#)–134
- Authorization models, [131](#)
- Availability of information, [7](#)
- AWS (Amazon Web Services), [96](#), [167](#)

B

- Backup data centers, [8](#)
- Backups, [95](#)–97
 - of AI systems, [236](#)
 - cloud, [96](#)
 - as corrective control, [56](#), [57](#)
 - as data control, [156](#)
 - differential, [96](#)–97
 - disk, [95](#)
 - disk mirroring, [87](#)
 - disk-to-disk, [95](#)
 - firewalls as, [84](#), [85](#)
 - full, [96](#)–97
 - incremental, [96](#)–97
 - offline storage of, [98](#)
 - personnel for, [207](#)

RAID vs., [88](#)
snapshot, [96](#), [116](#), [218](#)
tape, [40](#), [95](#)
and testing, [99](#)
Baselining, [215](#), [218](#), [237](#)
BCP (Business Continuity Planning), [81](#)–90
cloud-centric environments, [83](#)
COOP, [81](#)
fault tolerance in, [85](#)
high availability, [85](#)
impact analysis in, [82](#)
networking components of, [88](#)
power supplies in, [86](#)
redundancy in, [84](#), [86](#), [89](#)
single point of failure analysis in, [83](#)–84
Best practices, [48](#), [51](#), [104](#), [108](#), [115](#), [146](#), [149](#)
BIA (Business Impact Analysis), [82](#)
Biometric authentication, [8](#), [18](#), [19](#), [20f](#)
Blue Teaming, [226](#), [231](#)
Bluetooth, [140](#)
Box, [167](#)
Business alignment, [75](#)–76

C

Captive portals, [148](#)–149, [148f](#), [149f](#)

CCMP, [149](#)

CCPA (California Consumer Privacy Act), [54](#)

Change management, [218](#), [219](#)–221, [220f](#), [223](#)

ChatGPT, [234](#)

Checksums, [235](#)

CIA (Confidentiality, Integrity, and Availability) triad, [3](#)–11, [4f](#), [81](#), [234](#)

Ciphertext, [174](#)–175, [180](#)

CISA (Cybersecurity and Infrastructure Security Agency), [196](#)

Claude, [234](#)

Clearing (of data), [186](#)–187

Cloud backups, [96](#)

Cloud-centric environments, [83](#)

Cloud computing:

as asset, [217](#)

In data life cycle, [186](#)

definition of, [163](#)–164

hybrid, [166](#)

IaaS in, [170](#)

PaaS in, [167](#)

private, [166](#), [170](#)

public, [167](#), [170](#)

SaaS in, [167](#)

service models of, [166](#)–167

vendors for, [96](#)

Cloud security, [163](#)–170

Clustered web servers, [83](#), [83f](#)

COBIT, [77](#), [80](#)

Code of Professional Conduct, [66](#)

Cold sites, [98](#), [102](#)

Collection and creation, [29](#), [35](#)

Command and Control (C2), [199](#)

Community cloud deployment, [166](#), [170](#)

Compliance tracking, [77](#)–78

Confidentiality threats, [4](#)–6, [11](#), [39](#), [157](#), [173](#), [180](#)

Configuration management, [218](#)–219

Configuration vulnerabilities, [219](#)

Consensus and social proof, [105](#), [112](#)

Continuous improvement, [66](#), [76](#), [209](#)

Continuous monitoring and validation, [157](#)

Control risk, [43](#), [43^f](#)

COOP (Continuity Of Operations Planning), [81](#), [84](#)

Corrective controls, [56](#), [59](#)

Counterintelligence programs, [212](#)

Cross-cut shredding, [187](#)

Cross-site scripting (XSS), [227](#)

Cryptography, [173](#)–175. *See also* [Encryption](#)

asymmetric, [177](#)

diversifying, [89](#)

elliptic curve, [180](#)

hash functions in, [178](#)

post-quantum, [181](#)

purging, [187](#)

quantum-resistant, [180](#)–181
weaknesses in, [219](#)
CSF (NIST Cybersecurity Framework), [48](#)–49, [54](#), [77](#), [117](#), [120](#)
CSPs (Cloud Service Providers), [83](#)
CTI (Cyber Threat Intelligence), [196](#)
Cyber Kill Chain, [198](#)–199, [199f](#)
Cybersecurity controls, [55](#)–59
Cybersecurity metrics, [113](#)–118

D

DAC (Discretionary Access Control), [133](#), [136](#)
Dashboards, [116](#), [120](#), [238](#)
Data anonymization, [188](#)
Data at rest, [178](#), [180](#), [183](#)
Data centers:
 damage to, [82](#)–83
 for disaster recovery, [97](#)–98
 governance over, [48](#)
 multipath networking approach for, [88](#)
 power supplies for, [86](#)
 risks to, [41](#)–42
 threats to, [93](#)
Data classification, [189](#)–190
Data controls, [156](#)
Data handling policies, [185](#)–192, [205](#), [207](#)

Data in transit, [178](#), [180](#), [183](#)

Data integrity, [32](#), [235](#), [239](#)

Data leakage, [236](#), [238](#), [241](#)

Data life cycle, [185](#)–188, [185f](#)

Data-Link layer, [142](#)

Data masking, [188](#)–189

Data sanitization, [187](#)

DBMS (Database Management System), [203](#)

DCS (Distributed Control Systems), [151](#)

DDoS (Distributed Denial-of-Service) attacks, [7](#), [195](#)

Decryption keys, [149](#), [175](#), [178](#). *See also* [Encryption keys](#)

Default configurations, [219](#)

Defense in depth, [56](#), [130](#), [156](#), [241](#)

Deprovisioning, [126](#)–128, [130](#), [237](#). *See also* [Provisioning](#)

Destination addresses, [143](#)–144

Destroying (of data), [186](#)–188

Detective controls, [56](#), [59](#), [126](#)

DHCP (Dynamic Host Configuration Protocol), [144](#)

Diamond Model of Intrusion Analysis, [199](#)–200, [200f](#)

Differential backups, [96](#)–97

Digital access control, [14](#)

Directory services, [128](#)

Disclosure to third parties, [30](#)–31

Disk backups, [95](#)

Disk mirroring, [87](#), [88f](#)

Disk striping with parity, [87](#), [88f](#), [91](#)

Disk-to-disk backups, [95](#)
Diversity, [89](#)
DMZ (Demilitarized Zone), [152](#), [158](#)–159, [162](#)
DNS (Domain Name System), [145](#)–146, [158](#)
Dormant accounts, [126](#)
DoS (Denial-of-Service) attacks, [7](#), [209](#), [236](#)
Dotted quad notation, [142](#)–143
DR (Disaster Recovery), [93](#)–100, [236](#)
Dropbox, [167](#)
Dual power supplies, [86](#), [91](#)
Due care, [62](#)–63
Due diligence, [63](#), [76](#), [241](#)
Dynamic analysis, [227](#), [228](#)
Dynamic data masking, [188](#)

E

Eavesdropping, [5](#)–7, [26](#), [146](#), [149](#), [178](#), [219](#)
ECC (Elliptic Curve Cryptography), [180](#)
EDR (Endpoint Detection and Response), [156](#)
EEC (Elliptic Curve Cryptography), [180](#)
Emergency changes, [221](#)
Encryption. *See also* [Cryptography](#)
asymmetric, [177](#)
in compliance policies, [51](#)
in data classification, [190](#)

data masking vs., [189](#)
as due care, [63](#)
for electronic eavesdropping, [6](#)
and expectation of privacy, [26](#), [31](#)
full-disk, [5](#), [173](#), [178](#), [183](#)
for MitM attacks, [6](#)
for non-repudiation, [8](#)
plaintext, [174](#)–175, [174f](#), [180](#)
in ransomware, [56](#)
RSA, [179](#)
as security control, [57](#)
symmetric, [175](#)–177, [176f](#)
as technical control, [59](#)
uses of, [178](#)
Encryption keys, [147](#), [149](#), [174](#), [176](#). *See also* [Decryption keys](#)
Endpoint systems, [158](#)
Enterprise authentication, [148](#)
EOL (End of Life) software, [217](#)–218, [223](#)
EOS (End of Support), [217](#)–218
Ethical conduct, [61](#)–69
in AI, [235](#)
Exam Guidance for Artificial Intelligence (ISC2), [234](#)
Exploitation, [196](#), [199](#)
External risks, [38](#)

F

Familiarity and liking, [106](#), [112](#)

FDE (Full-Disk Encryption), [5](#), [173](#), [178](#), [183](#)

File integrity monitoring, [235](#)

Firewall zones, [158](#)–159

Firewalls, [157](#)–159

additional, [46](#)

AI-powered, [237](#), [238](#)

alerts from, [194](#)

in clustered web servers, [83](#)–84, [83f](#)

data from, [210](#)

for DoS attacks, [7](#)

as due care, [63](#)

high availability, [84](#), [84f](#), [85](#)

in IoT, [151](#)

in IP addressing, [143](#)

logging, [116](#), [194](#)

network border, [158f](#)

as perimeter control, [156](#)

as preventative control, [56](#)

risk management by, [39](#), [43](#)

routers vs., [162](#)

in VPNs, [150](#)

FT (Fault Tolerance), [85](#)–86

Full backups, [96](#)–97

Full interruption tests, [100](#)

G

GDPR (General Data Protection Regulation), [48](#), [54](#), [74](#), [80](#)

GLBA (Gramm-Leach-Bliley Act), [48](#)

Google cloud, [167](#)

Google Gemini, [234](#)

Google workspace, [167](#)

Governance, [47](#)–54

GRC (Governance, Risk, and Compliance) program, [73](#)–80, [189](#), [235](#)–236

H

HA (High Availability), [85](#), [91](#)

Hacktivists, [195](#)

Hallucinations, [235](#)

Hash functions, [178](#)–180

High availability firewalls, [84](#), [84f](#)

HIPAA (Health Insurance Portability and Accountability Act), [26](#), [48](#), [54](#), [190](#)

HNDL (Harvest Now, Decrypt Later) attacks, [180](#)–181

Horizontal scaling, [164](#)

Hot sites, [97](#), [102](#)

HTTP protocol, [142](#), [178](#), [183](#)

HTTPS protocol, [158](#), [178](#), [183](#)

Hybrid cloud approach, [166](#)

I

IaaS (Infrastructure as a Service), [167](#), [169](#), [170](#)

ICMP (Internet Control Message Protocol), [141](#)

ICS (Industrial Control Systems), [151](#)

Identification, [14](#), [15](#), [17](#), [23](#), [26](#), [31](#). *See also* [PII \(Personally Identifiable Information\)](#)

Identity attribution, [193](#), [203](#). *See also* [Accountability](#)

Identity life cycle management, [123](#)–128, [130](#)

Identity management frameworks, [127](#)–128

IDS (Intrusion Detection Systems), [238](#)

Impersonation, [4](#), [6](#)–7, [11](#), [20](#), [107](#), [229](#), [231](#)

Impossible travel, [237](#), [241](#)

Incident detection, [210](#), [236](#)

Incident notification and escalation procedures, [209](#)

Incident reports, [120](#), [211](#)

Incident use cases, [195](#)

Incineration, [188](#)

Incremental backups, [96](#)–97

Inference, [233](#)

Inherent risk, [43](#), [43f](#)

Initial response, [94](#)

Insider threats, [195](#)

Intellectual property, [38](#), [39](#)

Internal risks, [38](#)

Intimidation, [105](#), [112](#)

Intrusion detection, [59](#), [156](#), [194](#), [195](#), [210](#), [237](#)

IoCs (Indicators of Compromise), [196](#)

IoT (Internet of Things), [125](#), [151](#)
IP addresses, [140](#), [142](#)–146, [154](#), [193](#), [196](#)
IP (Internet Protocol), [140](#)–141
IPv4, [142](#)–145
IPv6, [143](#)
IR (Incident Response), [205](#)–214, [230](#)
as administrative control, [156](#)
AI for, [236](#)
in Code of Professional Conduct, [66](#)
communication plans for, [209](#)–210
exercises for, [208](#)–209
and identification, [210](#)–212
KPIs for, [114](#)
metrics for, [116](#)
procedures for, [51](#)
teams for, [207](#)–208
IRPs (Incident Response Plans), [205](#)–210, [236](#)
ISACA, [77](#)
ISACs (Information Sharing and Analysis Centers), [196](#)
ISC2, [234](#)
ISC2 Ethics Committee, [65](#)
ISMS (Information Security Management Systems), [51](#), [77](#)
ISO 27001, [51](#), [77](#)
ISO 31000, [77](#), [80](#)
ISO (International Organization for Standardization), [51](#)
Isolation, [165](#), [194](#)

IT contingency planning, [84](#)

K

Keccak algorithm, [180](#)

KPIs (Key Performance Indicators), [114](#)–116, [120](#)

KRIs (Key Risk Indicators), [115](#)–116

L

Labeling, [190](#)

LANs (Local Area Networks), [139](#)–140

LastPass, [18f](#)

Least privilege, [6](#), [125](#)–126, [130](#), [131](#)–132, [157](#), [234](#). *See also* [PoLP \(Principle of Least Privilege\)](#)

Legacy systems, [38](#)–39

LLMs (Large Language Models), [234](#), [235](#), [238](#)

Load balancing, [85](#), [91](#)

Lockheed Martin, [198](#)

Log monitoring, [194](#)

Logging, [193](#). *See also* [Access logging](#)

Logical access control, [31](#), [131](#)–136. *See also* [Access control](#)

Loopback addresses, [144](#)–145

LTO (Linear Tape-Open) tapes, [95](#), [96f](#)

M

MAC (Mandatory Access Control), [132](#)–133, [136](#)

macOS, [146](#), [147f](#)

Malware, [40](#), [106](#), [195](#), [196](#), [199](#), [211](#), [235](#). *See also* [Antimalware software](#)

Mantraps, [229](#)

MD5 (Message Digest 5) hash function, [179](#)–180

Media sanitization, [186](#)–187

Message digest, [179](#)

MFA (Multi-Factor Authentication), [20](#)–21, [38](#), [107](#), [128](#), [156](#), [237](#)

Microsegmentation, [157](#), [160](#). *See also* [Network segmentation](#); [VLANs \(virtual LANs\)](#)

Microsoft 365, [167](#)

Microsoft Active Directory (AD), [128](#)

Microsoft Azure, [96](#), [167](#)

MitM (Man-in-the-Middle) attacks, [6](#)–7

Mitnick, Kevin, [105](#)

MITRE ATT&CK, [196](#)–197, [197f](#), [228](#)

ML (Machine Learning), [233](#), [235](#)–238, [241](#)

Model drift, [236](#), [251](#)

Model poisoning, [235](#)–236, [241](#)

Monitoring and enforcement, [32](#)

MTTD (Mean Time To Detect), [114](#)

MTTR (Mean Time To Respond), [102](#)

Multicast addresses, [145](#)

Multiparty risks, [38](#)

Multipath networking, [88](#)

Multitenancy model, [165](#)

N

- NAS (Network-Attached Storage), [95](#)
- Nation-state actors, [195](#)
- Network border firewalls, [158f](#)
- Network controls, [156](#)
- Network layer, [142](#)
- Network ports, [157](#)–158
- Network security, [139](#)–154
 - embedded systems in, [150](#)–151
 - for IoT devices, [151](#)
 - IP addressing for, [142](#)–146
 - TCP/IP in, [140](#)–142
 - types of, [139](#)–140
 - VPNs for, [150](#)
 - for Wi-Fi networks, [146](#)–150
- Network security architecture, [155](#)–162
- Network segmentation, [152](#), [156](#), [157](#)–160, [218](#), [238](#). *See also* [Microsegmentation](#); [VLANs \(Virtual LANs\)](#)
- Networking components, [88](#)
- NFC (Near-Field Communication) technology, [140](#)
- NIC (Network Interface Cards) teaming, [89](#)
- NIDS (Network Intrusion Detection System), [203](#). *See also* [Intrusion detection](#)
- NIST Computer Security Incident Handling Guide*, [206](#)
- NIST Cybersecurity Framework (CSF), [48](#)–49, [54](#), [77](#), [117](#), [120](#)
- NIST Risk Management Framework (RMF), [49](#)–50

O

Offsite storage, [98](#)

One-time passcodes, [20](#)

OpenAI, [234](#)

Operational efficiency, [76](#)

Operational intelligence, [196](#)

Organizational culture, [103](#)–104

Organized cybercriminals, [195](#)

Orphaned accounts, [126](#)

OSI (Open Systems Interconnection) model, [141](#)–142, [142f](#)

OSINT (Open-Source Intelligence), [196](#)

P

PaaS (Platform as a Service), [167](#)–168, [170](#)

PAM (Privileged Access Management), [128](#)

PANs (Personal Area Networks), [140](#)

Parallel tests, [99](#), [100](#)

Password managers, [17](#)–18

Password policies, [16](#)–18

Patch compliance rate, [114](#)

Patch management, [219](#)

PCI DSS (Payment Card Industry Data Security Standard), [48](#), [54](#), [74](#), [80](#).
See also [Security standards](#)

PdUs (Power Distribution Units), [86](#)

Perimeter controls, [156](#)

PHI (Protected Health Information), [26](#), [35](#), [54](#)

Phishing, [4–5](#), [20](#), [106–107](#), [195](#), [196](#), [199](#), [228](#), [231](#), [234](#), [235](#)

Phishing click rate, [114](#)

Physical controls, [57](#), [59](#), [156](#)

Physical layer, [142](#)

Physical penetration testing, [228–229](#)

Physical security, [31](#), [208](#), [226](#), [228](#)

PII (Personally Identifiable Information), [26](#), [54](#), [189](#). *See also* [Identification](#)

PINs (Personal Identification Numbers), [5](#), [19](#), [21](#), [23](#)

Plaintext encryption, [174–175](#), [174f](#), [180](#)

PLCs (Programmable Logic Controllers), [151](#)

PMF (Privacy Management Framework), [27–32](#), [34–35](#). *See also* [Privacy](#)

Policy enforcement, [157](#)

Policy management, [78](#)

PoLP (Principle of Least Privilege), [131](#), [132](#), [237](#), [241](#). *See also* [Least privilege](#)

Potential impacts, [37](#), [74](#), [82](#), [82t](#), [157](#), [228](#)

Power supplies, [86](#), [86f](#), [88](#), [91](#)

PQC (Post-Quantum Cryptography), [181](#)

Presentation layer, [142](#)

Prevention systems, [57](#), [82–83](#), [156](#), [210](#)

Preventive controls, [56–57](#), [59](#)

Principals, [63–65](#), [69](#)

Prioritized risks, [82t](#)

Privacy, [25–35](#). *See also* [PMF \(Privacy Management Framework\)](#)
access to information, [30](#)

codes for, [66](#), [80](#)
disclosure of information, [30](#)–31
laws for, [210](#)
reasonable expectation of, [26](#), [27f](#), [35](#)
risk assessments for, [28](#)
risks of, [49](#)
security for, [31](#)
types of information of, [26](#)
use, retention, and disposal of information for, [29](#)–30
Privacy policies, [15](#), [28](#)–29, [32](#)
Privacy screens, [5](#)
Private cloud, [165](#), [166](#), [170](#)
Private keys, [177](#)
Prompt injection, [236](#)
Provisioning, [124](#)–126, [128](#), [130](#), [237](#). *See also* [Deprovisioning](#)
Pseudonymization, [188](#)–189
PSKs (Pre-Shared Keys), [147](#)–148
Public cloud, [164](#)–168, [170](#)
Public keys, [177](#)–178, [181](#)
Pulping, [188](#)
Purging (of data), [187](#)
Purple Teaming, [226](#), [231](#)

Q

QoS (Quality of Service), [143](#)

Qualitative risk assessment, [41](#), [41f](#)
Quantitative risk assessment, [41](#)
Quantum computing, [180](#)
Quantum-resistant cryptography, [180](#)–181

R

RaaS (Ransomware-as-a-Service), [195](#)
RAID (Redundant Arrays of Independent Disks), [87](#), [88f](#), [90](#), [91](#)
Ransomware, [38](#), [56](#), [93](#), [195](#), [209](#)
RBAC (role-based access control), [125](#), [134](#), [136](#)
Read-throughs, [99](#), [100](#)
Reasonable expectation of privacy, [26](#), [27f](#), [35](#)
Reconnaissance, [199](#), [233](#)–234
Red teaming, [209](#), [225](#)–226, [231](#)
Redundancy, [97](#)
Redundant network links, [84f](#)
Regulatory compliance, [75](#)–76
Replay attacks, [7](#), [11](#)
Reports, [117](#), [120](#)
Residual risk, [43](#), [43f](#)
Risk, categories of, [37](#)–39
Risk acceptance, [42](#)–43, [46](#)
Risk assessment, [236](#)
in AI, [236](#)
in BIA, [82](#)

as due diligence, [63](#)
in GRC approach, [75](#)–76
in offsite storage, [98](#)
in privacy management, [28](#), [32](#)
qualitative, [41](#), [41f](#)
quantitative, [41](#)
techniques, [41](#)
Risk avoidance, [42](#), [46](#)
Risk likelihood, [40](#)
Risk management, [37](#)–46. *See also* [RMF \(NIST Risk Management Framework\)](#)
as administrative control, [57](#)
definition of, [37](#)
goal of, [43](#)
in GRF framework, [74](#), [75](#)–76
life cycle of, [38f](#)
in PCI DSS framework, [79](#)
understanding, [74](#)
Risk mitigation, [42](#), [46](#)
Risk profiles, [43](#), [51](#)
Risk registers, [77](#)
Risk tolerance, [43](#)
Risk transference, [42](#), [46](#)
RMF (NIST Risk Management Framework), [49](#)–50, [50f](#), [77](#). *See also* [Risk management](#)
Roles definition, [125](#)
RPO (Recovery Point Objective), [94](#)–95, [102](#)

RSA (Rivest, Shamir, and Adleman) algorithm, [177](#), [179](#), [180](#)

RSL (Recovery Service Level), [94](#)–95, [102](#)

RTO (Recovery Time Objective), [94](#)–95, [102](#)

Runtime errors, [227](#)

S

SaaS (Software as a Service), [38](#), [167](#), [168](#), [170](#)

SAE (Simultaneous Authentication of Equals), [149](#)

SANs (Storage Area Networks), [95](#)

SCADA (Supervisory Control and Data Acquisition) systems, [151](#)

Scarcity, [106](#), [112](#)

Scorecards, [116](#)–117, [120](#)

Script kiddies, [195](#)

Secure communications, [157](#), [210](#)

Security awareness, [63](#), [103](#)–112, [109f](#)

Security awareness training, [57](#), [114](#), [229](#), [239](#)

Security baselines, [215](#)–216, [218](#). *See also* [Baselining](#)

Security classifications, [189](#)

Security culture, [104](#), [229](#), [235](#)

Security event triage, [194](#)–195, [236](#)

Security for privacy, [31](#)

Security guards, [14](#), [57](#), [229](#)

Security leadership, [104](#), [113](#), [116](#)

Security operations, [193](#)–203

cyber threat intelligence for, [196](#)

event triage, [194](#)–195

log monitoring, [194](#)

logging, [193](#)

for threat actors, [195](#)

threat frameworks for, [196](#)–200

Security policies, [6](#), [50](#)–51, [54](#), [63](#), [80](#), [104](#), [156](#)–158, [160](#)

Security procedures, [51](#)

Security readiness testing, [225](#)–226, [231](#)

Security settings, [146](#), [218](#), [219](#)

Security standards, [51](#), [215](#)–216. *See also* [PCI DSS \(Payment Card Industry Data Security Standard\)](#)

Security testing, [225](#)–231

Security training, [108](#)–109, [115](#)

Security vestibules, [229](#)

Server configuration errors, [227](#)

Session layer, [142](#)

SHA (Secure Hash Algorithm), [179](#)–180

Shared responsibility model, [165](#), [167](#)–168

SIEM (Security Information and Event Management), [194](#)–196, [203](#), [211](#), [238](#)

Simulations, [99](#), [100](#), [209](#)

Site resiliency, [98](#)

Smart cards, [20](#)–21, [23](#)

Smishing (SMS phishing), [107](#), [228](#)

SMTP, [142](#), [158](#)

Snapshots, [96](#), [116](#), [218](#)

SOCs (Security Operations Centers), [116](#), [120](#), [194](#)–195, [226](#)

SoD (Separation of Duties), [132](#)
Software license compliance, [38](#)–39
Source addresses, [143](#)–144
Spam filters, [235](#)
Spear phishing, [5](#), [107](#), [228](#)
SPOF (Single Point of Failure) analysis, [83](#), [84](#)
SQL injections, [46](#), [194](#), [227](#)
SSH, [142](#), [158](#)
SSID (Service Set Identifier), [146](#), [147f](#)
SSO (Single Sign-On), [128](#)
Staffing, [94](#)
Stakeholder confidence, [78](#)
Standing, [65](#), [68](#), [69](#)
Static analysis, [227](#), [228](#)
Static data masking, [188](#)
Static IP addresses, [144](#)
Strategic intelligence, [196](#), [212](#)
Strong dynamic identity verification, [156](#)
Symmetric encryption, [175](#)–177, [176f](#)

T

Tabletop exercises, [99](#), [209](#)
Tactical intelligence, [196](#)
Tailgating, [229](#), [231](#)
Tape backups, [95](#)

TCP/IP (Transmission Control Protocol/Internet Protocol), [140](#)–142

TCP (Transmission Control Protocol), [141](#)

Technical controls, [16](#), [57](#), [59](#), [103](#), [107](#), [110](#)

Threat actors, [106](#), [195](#), [233](#)–234

Threat frameworks, [196](#)–200

Threat modeling, [228](#)

TKIP (Temporal Key Integrity Protocol), [149](#)

TLS (Transport Layer Security), [7](#)

Tokenization, [5](#), [20](#), [188](#)

Traceability, [193](#), [203](#)

Training data, [233](#), [235](#)–236, [238](#), [241](#)

Transparency, [62](#), [66](#), [235](#)

Transport layer, [142](#)

TTPs (Tactics, Techniques, and Procedures), [196](#), [226](#)

U

UDP (User Datagram Protocol), [141](#), [142](#)

Unauthenticated scans, [227](#)

Unified reporting, [76](#)

Update management, [219](#)

UPSs (Uninterruptible Power Supplies), [86](#), [87](#)^f

Urgency, [106](#), [107](#)

URL (Uniform Resource Locator) addresses, [145](#)

Use, retention, and disposal (of information), [29](#)–30

V

Version control, [225](#)

Vertical scaling, [164](#)

Vishing, [107](#), [228](#), [231](#)

VLANs (virtual LANs), [156](#), [159](#)–160, [160f](#). *See also* [Microsegmentation](#); [Network segmentation](#)

VPNs (Virtual Private Networks), [22](#), [150](#), [166](#), [178](#), [183](#)

Vulnerability remediation, [114](#)

Vulnerability scanning, [63](#), [120](#), [211](#), [227](#)

W

Walk-throughs, [99](#), [100](#)

WANs (Wide Area Networks), [139](#)

Warm sites, [98](#), [102](#)

Weaponization, [199](#)

Web-based applications, [83](#), [83f](#)

WEP (Wired Equivalent Privacy) protocol, [149](#), [154](#)

Whaling, [5](#), [107](#), [228](#)

WPA (Wi-Fi Protected Access), [149](#)

WPA2, [149](#), [154](#)

WPA3, [149](#), [154](#)

X

XSS (Cross-Site Scripting), [227](#)

Z

ZT (Zero Trust), [156](#)–157, [160](#), [238](#)

CCSM Certified in Cybersecurity Study Guide, Second Edition. Mike Chapple.

© 2026 John Wiley & Sons, Inc. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published 2026 by John Wiley & Sons, Inc. Companion Website: https://www.wiley.com/go/chapple/CC_CertCySG_2e

Get Certified!

Join One of CertMike's FREE Study Groups and Get the Training You Need on a Timeline You Can Manage.

Mike Chapple, Ph.D. offers **FREE ONLINE STUDY GROUPS** that complement this book and will help prepare you for your next technology certification.

Mike is a cybersecurity and IT certification expert, professor, and author with over 25 years of experience in the field. Through his video courses, books, and study groups, he's helped millions of students earn professional IT certifications. He's written over 50 books, including the Official ISC2 CISSP Study Guide and created over 150 video courses helping students advance in their professional careers.



Entry Level Security Certifications

Security+
CC
SSCP



Advanced Security Certifications

CISSP PenTest+
CISM CCSP
CySA+



Privacy Certifications

CIPP/US
CIPM



IT and Data Certifications

Tech+
Data+
DataSys+

Earn Your Next Certification!
Visit CertMike.com to learn more.



Online Test Bank

To help you study for your Certified in Cybersecurity certification exam, register to gain one year of FREE access after activation to the online interactive test bank—including with your purchase of this book!

To access our learning environment, simply visit www.wiley.com/go/sybextestprep, follow the instructions to register your book, and instantly gain one year of FREE access after activation to:

- ▶ A complete practice exam
- ▶ Electronic flash cards
- ▶ A searchable glossary
- ▶ Exam essentials audio reviews

Wiley End User License Agreement

Go to www.wiley.com/go/eula to access Wiley's ebook EULA.