CRISC

Certified in Risk and Information Systems Control

STUDY GUIDE

2025 JOB PRACTICE

Includes interactive online learning environment and study tools:

2 custom practice exams

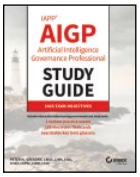
100 electronic flashcards

Searchable key term glossary

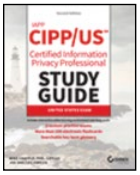
PETER H. GREGORY
CRISC, CISM, CISA, CISSP, CDPSE, CIPM, CCSK

SYBEX
A Wiley Brand

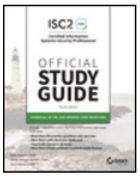
Other Information Security Study Guides by Sybex



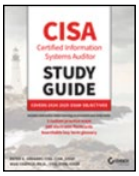
IAPP AIGP Artificial Intelligence Governance Professional Study Guide — ISBN: 978-1-394-36394-0, January 2026



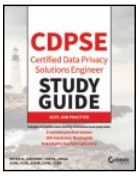
IAPP CIPP / US Certified Information Privacy Professional Study Guide, 2nd Edition — ISBN 978-1-394-28490-0, January 2025



ISC2 CISSP Certified Information Systems Security Professional Official Study Guide, 10th Edition — ISBN 978-1-394-25469-9, June 2024



CISA Certified Information Systems Auditor Study Guide Covering 2024–2029 Exam Objectives — ISBN 978-1-394-28838-0, December 2024



CDPSE Certified Data Privacy Solutions Engineer Study Guide — ISBN 978-1-394-36380-3, *coming September 2026*

CRISC

Certified in Risk and Information Systems Control

STUDY GUIDE



CRISC

Certified in Risk and Information Systems Control

STUDY GUIDE



PETER H. GREGORY
CRISC, CISM, CISA, CISSP, CDPSE, CIPM, CCSK

SYBEX®
A Wiley Brand

Copyright © 2026 by John Wiley & Sons, Inc. All rights, including for text and data mining, AI training, and similar technologies, are reserved. Some content was previously published in CRISC Certified in Risk and Information Systems Control All-In-One Exam Guide, Second Edition, by Bobby E. Rogers, Kimberley Dawn Dunkerley, And Peter H. Gregory (© 2022 McGraw-Hill).

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. CRISC is a trademark or registered trademark of Information Systems Audit and Control Association, Inc. 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. 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 website at www.wiley.com.

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

Paperback ISBN: 9781394373666
ePDF ISBN: 9781394373680
ePub ISBN: 9781394373673
oBook ISBN: 9781394373697

Cover Design: Wiley
Cover Image: © Jeremy Woodhouse/Getty Images

To cybersecurity and risk management professionals everywhere who are working to reduce cybersecurity-related risks in their organizations.

Acknowledgments

I am immensely grateful to James Minatel for affirming the need to have this book published on a tight timeline. My readers, including current and future risk managers, deserve nothing less.

Heartfelt thanks to Jim Minatel and Brad Jones for proficiently managing this project, facilitating rapid turnaround, and equipping me with the information and guidance we needed to produce the manuscript.

Many thanks to Maduramuthu Krishnaraj, Pete Gaughan, and Navin Vijayakumar for managing the editorial and production ends of the project, and to Kezia Endsley for copyediting the book and further improving readability. Like stage performers, they make hard work look easy, and I appreciate their skills.

I am grateful to Bobby Rogers and Dawn Dunkerley, co-authors of the *CRISC All-In-One Exam Guide* published by McGraw-Hill, for recruiting me as lead author of the second edition of the same book and publisher. And for this edition, published by Wiley Publishing, I am the sole author, but blessed by Bobby Rogers' continuing role with this manuscript, here in the role of technical editor. Bobby and I have worked on several book projects together, and I'm grateful for his deep expertise in information security and risk management.

Many thanks to my literary agent, Carole Jelen, for her diligent assistance during this and other projects. Sincere thanks to Rebecca Steele, my business manager and publicist, for her long-term vision and for keeping me on track.

Despite having written more than 50 books, I have difficulty putting into words my gratitude for my wife, Rebekah, for tolerating my frequent absences (in the home office) while I developed the manuscript. This project could not have been completed without her loyal and unfailing support and encouragement.

About the Author

Peter H. Gregory, CRISC, CISM®, CISA®, CDPSE™, CIPM®, CISSP®, CCSK™, is a 30-year career technologist and a security leader in a regional telecommunications company. He has been developing and managing information security programs since 2002 and leading the development and testing of secure IT environments since 1990. Peter has also spent many years as a software engineer and architect, a systems engineer, a network engineer, and a security engineer.

Peter is the author of more than 50 books about information security and technology, including *Solaris Security*, *CISM Certified Information Security Manager All-In-One Exam Guide*, and *CISA Certified Information Systems Auditor Study Guide*. He has spoken at numerous industry conferences, including RSA, Interop, ISC2 Congress, ISACA CACS, SecureWorld Expo, West Coast Security Forum, IP3, Source, Society for Information Management, the Washington Technology Industry Association, and InfraGard.

Peter serves on advisory boards for cybersecurity education programs at the University of Washington. He was the lead instructor for nine years in the University of Washington certificate program in applied cybersecurity, a former board member of the Washington State chapter of InfraGard, and a founding member of the Pacific CISO Forum. Peter is a 2008 graduate of the FBI Citizens Academy. He is a former member of the Forbes Technology Council.

Peter resides in Washington state with his family and can be found online at `www.peterhgregory.com`.

About the Technical Editor

Bobby E. Rogers is an information security engineer working as a contractor for Department of Defense agencies, helping to secure, certify, and accredit their information systems. His duties include information system security engineering, risk management, and certification and accreditation efforts. He retired after 21 years in the U.S. Air Force, serving as a network security engineer and instructor, and he has secured networks worldwide. Bobby has a master's degree in information assurance (IA). His many certifications include CRISC, CISSP-ISSEP, CEH™, and MCSE: Security, as well as the CompTIA A+™, Network+™, and Security+ certifications.

Contents at a Glance

Introduction *xxiii*

Assessment Test *xliii*

Answers to Assessment Questions *xlvii*

Chapter 1 Governance 1

Chapter 2 Risk Assessment 29

Chapter 3 Risk Response and Reporting 97

Chapter 4 Information Technology 163

Chapter 5 Information Security 223

Appendix Implementing and Managing a Risk Management Program 271

Index *295*

Contents

Introduction *xxiii*

Assessment Test *xliii*

Answers to Assessment Questions *xlvii*

Chapter 1 Governance 1

Organizational Governance 2
- Organizational Strategy, Goals, and Objectives 4
- Organizational Structure, Roles, and Responsibilities 4
- Organizational Culture 6
- Policies and Standards 7
- Business Processes 8
- Organizational Assets 9
  - Asset Identification 10
  - Sources of Asset Data 11
  - Collecting and Organizing Asset Data 12

Risk Governance 13
- ERM and Risk Management Frameworks 15
  - ISACA Risk IT Framework 15
  - NIST Risk Management Framework 15
  - COSO Enterprise Risk Management 16
  - ISO 31000 16
- Three Lines of Defense 17
  - Operational Management 17
  - Risk and Compliance 17
  - Internal Audit and Accountability 18
- Risk Profile 18
- Risk Appetite and Risk Tolerance 18
- Legal, Regulatory, and Contractual Requirements 20
- Professional Ethics of Risk Management 20

Summary 21
Exam Essentials 22
Review Questions 23
Answers to Review Questions 26

Chapter 2 Risk Assessment 29

IT Risk Identification 31
- Risk Events 34
  - The Vocabulary of Risk 34
  - Types of Risk Events 36

Threat Modeling and Threat Landscape 37
Threat Modeling 37
Threat Landscape 38
Vulnerability and Control Deficiency Analysis 40
Vulnerability Analysis 40
Control Deficiency Analysis 40
Evaluating Existing Control Effectiveness 41
Control Deficiencies and Gaps 43
Control Recommendations 44
Considering Risk and Control Ownership 45
Root Cause Analysis 45
Risk Scenario Development 47
Top-down Risk Scenario Development 48
Bottom-up Risk Scenario Development 48
Risk Scenario Development Considerations 49
IT Risk Analysis and Evaluation 49
Risk Assessment Concepts, Standards, and Frameworks 50
Types of Assessments 50
Scope of an Assessment 51
Collecting Evidence for an Assessment 51
Risk Assessment Standards and Frameworks 55
NIST Special Publication 800-30 55
OCTAVE Methodology 57
ISO/IEC Standards 59
ISACA's Risk IT Framework 59
Risk Ranking 62
Risk Ownership 62
Risk Register 63
Sources of Information for the Risk Register 67
Updating the Risk Register 67
Risk Analysis Methodologies 69
Asset Identification, Classification, and Valuation 69
Likelihood Determination and Impact Analysis 70
Quantitative and Qualitative Analysis Techniques 71
Other Analysis Techniques 75
Business Impact Analysis 78
Inventory Key Processes and Systems 79
Statements of Impact 80
Criticality Analysis 81
Recovery Objectives and Next Steps 83
Inherent and Residual Risk 83
Inherent Risk 83
Residual Risk 83

Miscellaneous Risk Considerations 85
Using AI to Identify and Analyze Risk 86
Role of AI in Risk Identification 86
AI in Risk Analysis and Evaluation 86
Integration with Frameworks and Standards 87
Benefits and Limitations 87
Governance and Human Oversight 88
Summary 88
Exam Essentials 90
Review Questions 91
Answers to Review Questions 94

Chapter 3 Risk Response and Reporting 97

Risk Response 99
Risk Treatment/Risk Response Options 99
Risk Mitigation 100
Risk Sharing 100
Risk Acceptance 101
Risk Avoidance 101
Evaluating Risk Response Options 102
Risk and Control Ownership 102
Third-party Risk Management 103
Risk Sharing Through Third Parties 104
Accountability and Responsibility in Third-party Relationships 104
Third-party Risk Factors 104
External Ratings and Continuous Third-party Monitoring 105
Third-party Threats 106
Third-party Vulnerabilities 106
Cloud Shared Responsibility Models 107
Third-party Risk Management 111
Business Process Outsourcing 111
Issues, Findings, and Exception Management 111
Change and Configuration Management 112
Exception Management 113
Workflow for Issues, Findings, and Exception Management 113
Management of Emerging Risk 116
Emerging Technologies 116
Emerging Risk Factors 117
Emerging Threats 118
Emerging Vulnerabilities 118
AI and Machine Learning as Emerging Risk Domains 119

Control Design and Implementation 120
Control Types, Standards, and Frameworks 120
Types of Controls 121
Control Standards and Frameworks 124
Control Design, Selection, and Analysis 129
Control Rationalization 129
Control Design 131
Control Selection 132
Control Analysis 133
Control Implementation 134
Documenting and Publishing Controls 134
Change and Configuration Management 135
Implementation Testing 135
Other Implementation Considerations 136
Control Testing and Effectiveness Evaluation 136
Continuous Monitoring and Assurance 137
Risk Monitoring and Reporting 139
Risk Treatment Plans 139
Data Collection, Aggregation, Analysis, and Validation 142
Risk and Control Monitoring Techniques 144
Risk and Control Reporting Techniques 144
Stakeholder-specific Reporting Needs 145
Standard Elements of Risk Reports 145
Incorporating Audit and Assessment Results 146
Creating Actionable and Decision-ready Reports 146
Communicating Risk Scenarios and Business Impact 146
Key Performance Indicators 147
Key Risk Indicators 149
Leading Indicators 150
Lagging Indicators 150
Using Leading and Lagging Indicators Together 150
Key Control Indicators 151
Summary 152
Exam Essentials 152
Review Questions 154
Answers to Review Questions 160

Chapter 4 Information Technology 163

Enterprise Architecture (EA) 164
EA Concepts and Structures 164
Enterprise Architecture Frameworks 166
The Open Group Architecture Framework 166
The Zachman Framework 168
EA Risk Considerations 168

IT Operations Management 170
Core Operations Functions 171
Infrastructure and Platform Operations 172
End-user Services and Device Management 172
IT Service Desk and Incident Management 172
Problem Management 172
Change Management 172
Configuration Management 173
Monitoring, Logging, and Observability 173
Backup, Recovery, and Data Protection 173
Identity and Access Operations 173
Vendor and Service Provider Coordination 174
Lifecycle and Asset Management 174
Security Management 174
Operations Governance and Reporting 174
Operational Risk 175
Process-related Risks 175
Technology-related Risks 175
People-related Risks 175
Security-related Operational Risks 175
Vendor and Third-party Risks 176
Environmental and Physical Risks 176
Governance and Compliance Risks 176
Infrastructure Risk 176
Networks 176
Operating Systems 177
Gateways 177
Software 177
Databases 178
Platforms 178
Cloud 179
Project, Program, and Portfolio Management 179
Project Drivers and Risks 180
Scope Risk 180
Schedule Risk 180
Cost Risk 180
Project Threats 180
Project Vulnerabilities 181
Project Risk Monitoring 182
Other Project Risks 182
Modern Development Models and Risks 183
Agile 183
Scrum 183
DevOps/DevSecOps Pipelines 183
CI/CD, Infrastructure-as-Code 184

Disaster Recovery Management 184
DR Concepts and Planning Activities 187
DR Objectives and Scope 187
Recovery Time Objective (RTO) and Recovery Point Objective (RPO) 187
Business Impact Analysis Integration 188
Role of DRM in Business Continuity 188
Risk Assessment and Scenario Identification 188
Plan Documentation 189
DR Training and Preparedness 189
Disaster Declaration, Activation Procedures 189
DR Strategies 190
DR Strategy Development 190
Backup-based Recovery 190
Redundant or High-availability Architectures 191
Replication and Mirroring 191
Alternate Processing Sites 191
Rebuild-from-code and Automation 191
Prioritized Recovery Sequencing 192
DR Plan Testing 192
Types of DR Testing 192
DR Testing Metrics 193
Data Lifecycle Management 194
Data Governance 194
Data Ownership and Stewardship 195
Data Classification 195
The Data Lifecycle 197
Systems Development Lifecycle 198
SDLC Phases 198
Planning 198
Requirements 199
Design 199
Development 200
Testing 200
Implementation 200
Operations and Maintenance 201
Post-implementation Review 201
Retirement 201
The Roots of SDLC 202
Shifting Left 202
Tying Requirements to Controls and Features 203
SDLC Risks 203
Risk Factors 203
Threats 204
Vulnerabilities 204

Secure Development Practices 205
Security Requirements and Design Controls 205
Threat Modeling 205
Code Quality and Secure Coding Standards 205
Static Application Security Testing 206
Dynamic Application Security Testing 206
Software Composition Analysis 206
Code Review and Peer Validation 206
Container and Artifact Security 206
Secrets and Credential Management 207
Logging, Monitoring, and Telemetry Requirements 207
Continuous Improvement and Governance 207
Emerging Technologies 207
Risk Factors 207
Security 208
System Maintenance 208
Threats 208
Vulnerabilities 209
Cloud and Virtualization 209
Virtualization Concepts and Risks 210
Cloud Service Models (IaaS, PaaS, SaaS, XaaS) 210
Shared Responsibility Model 210
Identity and Access in Cloud Environments 211
Virtual Networks and Segmentation 211
Workload Protection and Configuration Management 211
Data Protection and Storage Risk 211
Provider Dependency and Resilience 211
Modern Architectures 212
AI Models 212
Quantum Computing 212
Containerization 212
Kubernetes and Orchestration 213
Microservices Architectures 213
API-first Architectures and API Security 213
Zero-trust Architecture (ZTA) 214
Edge Computing 214
IT/OT Convergence 214
Serverless and Function-based Computing 214
Summary 215
Exam Essentials 215
Review Questions 217
Answers to Review Questions 221

Chapter 5 Information Security 223

Information Security Concepts, Frameworks, and Standards 224
Confidentiality, Integrity, and Availability 225
Frameworks 226
NIST RMF 226
NIST CSF 228
ISO/IEC 27001/27002/27701/31000 229
COBIT 2019 230
The Risk IT Framework (ISACA) 232
Policies 234
Information Security Policy 234
Privacy Policy 234
Acceptable Use Policy (AUP) 235
Standards 236
Controls 236
Practices 237
Due Care, Due Diligence, and Due Process 238
Network and Systems Security 238
Vulnerability Management 239
Security Incident Response 239
Third-party Risk Management 241
Access Management 241
Security Architecture 247
Information Security Awareness Training 247
Designing Training Programs 248
Audience-specific Content 248
Behavioral Reinforcement 249
Leadership-level Risk Training 249
Training Delivery and Measurement 249
Business Continuity Management 250
BCM Foundations 250
Business Impact Analysis 251
Recovery Objectives 251
Resilience Engineering 252
Data Privacy and Data Protection Principles 254
Data Sensitivity and Classification 256
Data Protection 256
Privacy as a Risk Scenario 257
Privacy Regulatory Foundations 258
GDPR 258
EU AI Act 259
CCPA/CPRA 259
Sectoral Laws 260

Summary 261
Exam Essentials 262
Review Questions 264
Answers to Review Questions 268

Appendix Implementing and Managing a Risk Management Program 271

Today's Risk Landscape 272
Allegory of the Battlefield 274
What Is a Risk Management Program? 275
The Purpose of a Risk Management Program 276
Risk Management Program Foundations 277
The Risk Management Lifecycle 277
Risk Discovery 277
The Risk Register 278
Rating Risks—Qualitative 280
Rating Risks—Quantitative 282
Types of Risk Registers 282
Layers of Risk Registers 282
Communicating Between the Layers 283
Reviewing the Risk Register 283
Formal Risk Register Reviews 284
Identifying Candidates for Deeper Analysis 285
Performing Deeper Risk Analysis 285
The Risk Analysis Worksheet 286
Qualitative and Quantitative Risk 287
Selecting Risk Treatment Options 287
Considering Risk Appetite and Risk Tolerance 288
The Precedence of Prior Risk Treatment Decisions 288
Developing a Risk Treatment Recommendation 289
The Risk Treatment Decision 290
Risk Mitigation 290
Risk Acceptance 290
Risk Transference 291
Risk Avoidance 292
Publishing and Reporting 293
Risk Register 293
New Risks 293
Risk Review 293
Risk Treatment 294
Risk Mitigation Follow-up 294
Risk Renewals 294
Risk Prioritization 294

Index *295*

Introduction

Welcome to the Sybex Study Guide for ISACA's Certified in Risk and Information Systems Control (CRISC) exam! This book will help you study for, and successfully pass, one of ISACA's premier certification exams, the CRISC exam. This exam is designed to test your knowledge of a wide variety of topics related to risk management and information systems controls. The exam focuses on business and IT risk management in enterprise infrastructure, as well as on designing and implementing IT security controls to mitigate risks.

Every day, it seems, data is breached at some of the largest organizations in the world. Recently, we've seen breaches in the U.S. government, in the healthcare industry, and even at tech giants such as Microsoft, LinkedIn, and Facebook. Ransomware attacks alone are a plague on thousands of companies in nearly every country in the world. And to make matters worse, cybercriminals are using AI to develop their attacks more quickly and make them stealthier and more potent. No organization, no matter what size, is immune to the threat of data breaches or, for that matter, data theft or loss. However, effective risk management can reduce the likelihood of data breaches, strengthen business processes and IT infrastructure, and even improve their efficient use. Information system controls reduce the likelihood of adverse events having a significant impact on the organization and should be carefully planned and considered.

This book covers basic risk concepts, risk assessments, standards and frameworks, and information security control design and implementation. I also cover the essential concepts, terminology, and definitions that risk management practitioners and security professionals need to be effective in these areas. In the book's five main chapters, I cover all four top-level domains as well as the task and knowledge statements listed in the official ISACA exam objectives. Appendix discusses the practical steps a security leader can take to successfully implement a risk management program.

While you don't have to be an expert already in all the areas I discuss, having experience in some, such as risk concepts, helps. A good, broad background of experience and knowledge in information security will give you an advantage in your studies for this exam. Of course, you'll get a good background in all these subjects throughout the book.

Passing the CRISC exam not only places you in a class of professionals recognized for their experience and expertise in this field, but it also serves to quantify and validate your knowledge of advanced risk management and security topics. After passing this exam, you'll be able to show that not only are you qualified, but you are certified in these areas. This book is designed to help you get there.

Purpose of This Book

Let's get the obvious out of the way: this is a comprehensive study guide for the information security and risk management professional who needs a reliable reference for individual or group-led study for the CRISC certification. This book contains the information that CRISC

candidates are required to know. While this book is one source of information to help you prepare for the CRISC exam, it should not be thought of as the ultimate collection of *all* the knowledge and experience that ISACA expects qualified CRISC candidates to possess—no one publication covers all this information. The other thing you'll need, just as important as suitable study material in our minds, is experience. There's no substitute for practical, hands-on experience. You should make every effort to learn all aspects of the ISACA CRISC exam material I discuss in this book.

This book also serves as a reference for aspiring and practicing security and risk professionals and leaders. The content required to pass the CRISC exam is the same content that practicing security and risk professionals need to be familiar with in their day-to-day work. This book is a definitive CRISC exam study guide as well as a desk reference for those who have already earned their CRISC certification.

The pace of change in the information security and risk management industry and profession is high. Rather than cover every detail and nuance of laws, practices, standards, and techniques in information security and risk management, this book shows readers how to stay current in the profession. Indeed, the pace of change is one of many reasons that ISACA and other associations require continuous learning to retain one's certifications. It is crucial to understand key facts and practices in security and risk management, and to learn how to stay current as they continue to evolve. However, despite the high rate of change in information technology and information security, there's good news: the techniques for governance and risk management themselves change very slowly. The principles of governance, risk assessments, risk management, risk treatment, and controls are solid and time-proven. Much of this book is devoted to these practices.

This book is also invaluable for security and risk management professionals who are not in a leadership position. You will gain considerable insight into today's security and risk management challenges. This book is also helpful for IT, privacy, and business management professionals who work with risk management professionals and need a better understanding of what they are doing, how they do it, and why.

Finally, this book is an excellent guide for anyone exploring a career in information security and risk management. The study chapters explain all the relevant technologies, techniques, and processes used to manage a modern risk management program, which is helpful if you are wondering what the risk management profession is all about.

How to Use This Book

This book covers everything you'll need to know for ISACA's CRISC certification examination. Each chapter covers specific objectives and exam details, as defined by ISACA in its job practice areas. The chapters and their sections correspond precisely to the CRISC job practice that ISACA updates from time to time.

Each chapter has several components designed to effectively communicate the information you'll need for the exam.

The topics covered in each chapter are listed in the first section to help you map out your study.

- The **Summary** section of each chapter briefly explains the chapter, allowing you to easily understand what it covers.
- **Exam Essentials** focus on major exam topics and critical knowledge that you should take into the test. The Exam Essentials focus on the ISACA CRISC exam objectives.
- **Tips** are included in each chapter that offer great information on how concepts you'll study apply in a place that we like to call "the real world." Often, they give you a bit more information on a topic covered in the text that may appear in the exam.
- **Notes** may be included in a chapter as well. These are bits of information that are relevant to the discussion and that point out extra information.
- **Case Studies** are included to illustrate how an idea, concept, or standard can be put into practice.
- **Twenty or more practice questions** appear at the end of each chapter and are designed to allow you to attempt some exam questions on the topics covered in the domain.

Appendix is designed to help you understand the practical side of risk management, particularly for security leaders who need to develop or improve the risk management function in an organization. I have many years of experience; this is my gift to you so that you have the greatest chance to succeed.

About This Edition

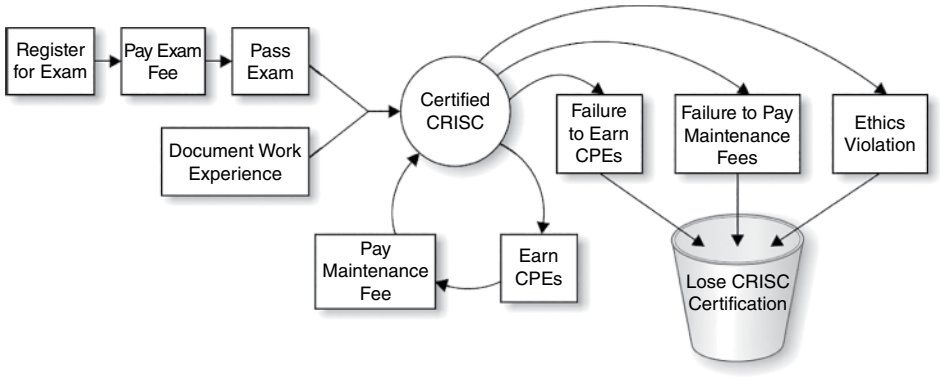
ISACA, like other certification organizations, periodically updates the job practice for CRISC and its other certifications. In 2024, I expected that ISACA would update the CRISC job practice in November 2025, which prompted me to prepare to write this edition reflecting these updates, as well as changes in related industry practices and developments.

ISACA is somewhat academic in its approach to the arrangement of the CRISC job practice areas. For this reason, I included an appendix that describes the development and operation of a risk management program from a practical perspective. This appendix helps to fill the gap between the CRISC job practice and what it really takes to succeed.

Becoming a CRISC Professional

To become a CRISC professional, you are required to pay the exam fee, pass the exam, prove that you have the necessary education and experience, and agree to uphold ethics and standards. To keep your CRISC certification, you are required to take at least 20 continuing education hours each year (120 hours in three years) and pay annual maintenance fees. This lifecycle is depicted in Figure 1.

FIGURE 1 The CRISC certification lifecycle.



The following list outlines the primary requirements for becoming certified:

- **Experience:** A CRISC candidate must submit verifiable evidence of at least three years of professional work experience in IT risk management and IS control. Experience must be verified and gained within the ten-year period preceding the application date for certification or within five years of passing the exam. No waiver options are available.
- **Ethics:** Candidates must commit to adhering to ISACA's Code of Professional Ethics, which guides the personal and professional conduct of those certified.
- **Exam:** Candidates must receive a passing score on the CRISC exam. A passing score is valid for up to five years, after which the passing score is void. A CRISC candidate who passes the exam has a maximum of five years to apply for CRISC certification; candidates who pass the exam but fail to act after five years will have to retake the exam if they want to become CRISC certified.
- **Education:** Those certified must adhere to the CRISC Continuing Professional Education Policy, which requires a minimum of 20 continuing professional education (CPE) hours each year, with a total requirement of 120 CPEs over the course of the certification period (three years).
- **Application:** After successfully passing the exam, meeting the experience requirements, and having read through the Code of Professional Ethics and Standards, a candidate is ready to apply for certification. An application must be received within five years of passing the exam.

Experience Requirements

To qualify for CRISC certification, you must have the equivalent of three years of total work experience in at least two of the CRISC domains. Additional details on the minimum certification requirements, substitution options, and various examples are discussed next.



Although it is not recommended, a CRISC candidate can take the exam before completing any work experience directly related to information security management. If the candidate passes the exam and meets the work experience requirements within five years of the exam date, and within ten years of the application for certification, the candidate is eligible for certification.

Direct Work Experience

You are required to have a minimum of three years of work experience in IT risk management and IS control, equivalent to 6,000 actual work hours, which must be related to at least two of the CRISC job practice areas, listed next. One of these two job practice areas must be in Governance or IT Risk Assessment.

- **Governance:** Establish and maintain organizational governance and risk governance programs and frameworks. Determine roles and responsibilities as well as develop policies, standards, business processes, risk appetite, and risk tolerance, aligned with business culture and compliance with legal, regulatory, and contractual requirements.
- **IT Risk Assessment:** Establish and maintain processes to identify and analyze risk, including vulnerability management, threat modeling, risk assessments, and business impact analysis. Develop and maintain a risk register.
- **Risk Response and Reporting:** Establish and maintain risk treatment and risk response, controls and control ownership, third-party risk management processes, and control testing and remediation. Establish and publish metrics, key performance indicators, key risk indicators, and key control indicators.
- **Information Technology and Security:** Understand the principles of information technology and security, including enterprise architecture, IT operations, project management, business continuity and disaster recovery planning, data lifecycle, systems development lifecycle, security awareness training, and privacy.

All work experience must be completed within the ten years before completing the certification application or within five years from initially passing the CRISC exam. You will need to complete a separate Verification of Work Experience form for each segment of experience.

Substitution of Experience

Unlike most other ISACA certifications, there are no available experience waivers or substitutions. Instead, you are required to have three or more years of experience in IT risk management and IS control.



I recommend you also read the CRISC certification qualifications on the ISACA website at `www.isaca.org/credentialing/crisc`. From time to time, ISACA changes the qualification rules, and I want you to have the most up-to-date information available.

ISACA Code of Professional Ethics

Becoming a CRISC professional means you agree to adhere to the ISACA Code of Professional Ethics, which is a formal document outlining those things you will do to ensure the utmost integrity and that best support and represent the organization and certification.

Specifically, the ISACA Code of Professional Ethics requires ISACA members and certification holders to do the following:

- Support the implementation of, and encourage compliance with, appropriate standards and procedures for the effective governance and management of enterprise information systems and technology, including audit, control, security, and risk management.
- Perform their duties with objectivity, due diligence, and professional care, in accordance with professional standards.
- Serve in the interest of stakeholders in a lawful manner, while maintaining high standards of conduct and character and not discrediting their profession or the association.
- Maintain the privacy and confidentiality of information obtained in the course of their activities unless disclosure is required by legal authority. Such information shall not be used for personal benefit or released to inappropriate parties.
- Maintain competency in their respective fields and agree to undertake only those activities they can reasonably expect to complete with the necessary skills, knowledge, and competence.
- Inform appropriate parties of the results of work performed, including the disclosure of all significant facts known to them that, if not disclosed, may distort the reporting of the results.
- Support the professional education of stakeholders in enhancing their understanding of the governance and management of enterprise information systems and technology, including audit, control, security, and risk management.

Failure to comply with the ISACA Code of Professional Ethics can result in an investigation into a member's or certification holder's conduct and, ultimately, disciplinary measures, including the forfeiture of their hard-won certification(s).

You can find the full text and terms of enforcement of the ISACA Code of Professional Ethics at `www.isaca.org/code-of-professional-ethics`.

The Certification Exam

The certification exam is offered year-round. You have several ways to register; however, I highly recommend you plan and register early, regardless of your chosen method.

In early 2026, the schedule of exam fees in U.S. dollars was:

- CRISC application fee: $50
- Regular registration: $575 member/$760 nonmember

The ISACA administers the test at testing centers while also offering remotely proctored exams for those who prefer remote testing. I discuss both options in this section.



I recommend you pay close attention to information on ISACA's website regarding testing logistics and locations.

Once registration is complete, you will immediately receive an email acknowledging your registration. Next, you will need to schedule your certification exam. The ISACA website will direct you to the certification registration page, where you will select a date, time, and (optionally) location to take your exam. When you confirm the date, time, and location for your exam, you will receive a confirmation via email. You will need the confirmation letter to enter the test location—make sure to keep it unmarked and in a safe place until test time.

Onsite Testing Center

When you arrive at the test site, you will be required to sign in, and you may be required to sign an agreement. Also, you will be required to turn in your smartphone, wallet or purse, and other personal items for safekeeping. The exam proctor will read aloud the rules you are required to follow while you take your exam. These rules will address matters such as breaks, drinking water, and snacks.

While you take your exam, you will be supervised by the proctor, and you will be monitored by video surveillance in the test center to make sure no one can cheat on the exam.

Remote Proctored Testing

If you have registered for a remote proctored exam, make sure you meet all the technical requirements. ISACA has published the "Remote Proctoring Guide," which includes all the necessary technical requirements and describes the step-by-step procedures for taking the exam. In the Exam Candidate Guide, go to the Online Remote Proctoring section, where you'll find a link to the Remote Proctoring Guide.

A remote proctored exam means you'll be taking the exam on your own computer in your residence or other location. You'll be in live contact with an exam proctor, and your webcam will be turned on throughout the exam so that the proctor can observe you while taking the exam to ensure you are not cheating with reference materials (books or online).

The proctor may ask you to show them the room where you are taking the exam to ensure you do not have reference materials or information anywhere in view.

To be eligible for a remote proctored exam, you must have a supported version of Windows or macOS, a current Google Chrome or other Chromium browser (such as Brave), a webcam with at least 640 × 480 resolution, a microphone, and a stable broadband Internet connection. You must have the ability to install the PSI Secure Browser and modify firewalls and other administrative tasks on the day of the exam (this requires you to have administrative privileges on the computer you are using, which might be a problem if you are using a company-issued computer that often restricts your ability to install software).

You'll be required to log in to your My ISACA account when your exam is scheduled. Next, you'll navigate to your certifications, find the exam you have scheduled, and launch the exam. You'll be directed to perform several tasks, including installing the secure browser and closing several other programs on your computer, including other web browsers and programs like Adobe Reader, Word, Excel, and any others that could include reference material.

You are not permitted to speak or perform gestures at any time during the exam. In short, you cannot be seen to perform any action that might be an indication of aid by an accomplice.

You'll be required to verify your ID by holding it near your webcam so that the proctor can see it to confirm that you are not having someone else take the exam for you. You will also be required to use your webcam to show the entire room to your proctor.

After all steps have been completed, the proctor will release the exam to you, and you may begin.

Exam Questions

Each registrant has four hours to take the multiple-choice question exam. There are 150 questions on the exam, representing the four job practice areas. Each question has four answer choices, of which you can select only one best answer. You can skip questions and return to them later, and you can also flag questions that you want to review later if time permits. While you are taking your exam, the time remaining will appear on the screen.

When you have completed the exam, you are directed to close it. At that time, the exam will display your pass or fail status, with a reminder that your score and passing status are subject to review.

You will be scored in each job practice area and then provided a final score. All scores are scaled. Scores range from 200 to 800; however, a final score of 450 is required to pass.

Exam questions are derived from a job practice analysis study conducted by ISACA. The areas selected represent those tasks performed in a CRISC's day-to-day activities and represent the background knowledge required to develop and manage an information security program. You can find more detailed descriptions of the task and knowledge statements at `https://www.isaca.org/credentialing/crisc/crisc-exam-content-outline`.

Exam Coverage

The CRISC exam is quite broad in its scope. The exam covers four job practice areas, as shown in Table 1.

TABLE 1 CRISC exam practice areas.

Domain	CRISC Job Practice Area	Percentage of Exam
1	Governance	26%
2	Risk Assessment	22%
3	Risk Response and Reporting	32%
4	Technology and Security	20%

Independent committees have been developed to determine the best questions, review exam results, and statistically analyze the results for continuous improvement. Should you come across a horrifically difficult or strange question, do not panic. This question may have been written for another purpose. A few questions on the exam are included for research and analysis purposes and will not be counted against your score. The exam contains no indications in this regard.

Preparing for the Exam

The CRISC certification requires a significant amount of knowledge and experience from the candidate. You need to develop a long-term study plan to pass the exam. The following sections offer some tips and are intended to help guide you to, through, and beyond exam day.

Before the Exam

Consider the following list of tips on tasks and resources for exam preparation. They are listed in sequential order.

- **Read the candidate's guide:** For information on the certification exam and requirements for the current year, see the ISACA Exam Candidates Information Guide. Go to `www.isaca.org/credentialing/crisc` and look for the Exam Candidate Guide download links.
- **Register:** If you are able, register early for any cost savings and solidify your commitment to moving forward with this professional achievement.

- **Schedule your exam:** Find a location, date, and time, and commit.
- **Become familiar with the CRISC job practice areas:** The job practice areas serve as the basis for the exam and requirements. Ensure your study materials align with the current list, shown at `www.isaca.org/credentialing/crisc`.
- **Know your best learning methods:** Everyone has preferred learning styles, whether self-study, a study group, an instructor-led course, or a boot camp. Try to set up a study program that leverages your strengths.
- **Self-assess:** Run through practice exam questions available in this book and online. ISACA may offer a free CRISC self-assessment at `www.isaca.org/credentialing/crisc`.
- **Study iteratively:** Depending on how much work experience in information security management you already have, I suggest you plan your study program to take at least two months, but as long as six months. During this time, periodically take practice exams and note your strengths and weaknesses. Once you have identified your weak areas, focus on those areas weekly by rereading the related sections in this book, retaking practice exams, and noting your progress.
- **Avoid cramming:** We've all seen the books on the shelves with titles that involve last-minute cramming. Just one look on the Internet reveals various websites that teach individuals how to cram for exams most effectively. Research sites claim that exam cramming can lead to susceptibility to colds and flu, sleep disruptions, overeating, and digestive problems. One thing is sure: many people find that good, steady study habits result in less stress and greater clarity and focus during the exam. Because of the complexity of this exam, I highly recommend the long-term, steady-study option. Study the job practice areas thoroughly. There are many study options. If time permits, investigate the many resources available to you.
- **Find a study group:** Many ISACA chapters and other organizations have formed specific study groups or offer less expensive exam review courses. Contact your local chapter to see whether these options are available to you. In addition, be sure to keep your eye on the ISACA website. Also, search online or use your local network to find out whether there are other local study groups and other helpful resources.
- **Recheck your confirmation letter:** Do not write on or lose your confirmation letter. Put it in a safe place and note what time you will need to arrive at the site. Note this on your calendar. Confirm that the location is the one you selected and is located near you.
- **Check logistics:** Check the candidate's guide and your confirmation letter for the exact time you are required to report to the test site (or log in from home if you registered for a remote proctored exam). A few days before the exam, check the site—become familiar with the location and tricks to getting there. If you are taking public transportation, be sure to check the schedule for the day of the exam. If your CRISC exam is scheduled on a Saturday, public transportation schedules may differ from weekday schedules. If you are driving, know the route and where to park.

- **Pack:** Place your confirmation letter and a photo ID in a safe place, ready to go. Your ID must be a current, government-issued photo ID that matches the name on the confirmation letter and must not be handwritten. Examples of acceptable forms of ID are passports, driver's licenses, state IDs, green cards, and national IDs. Ensure you leave food, drinks, laptops, cell phones, and other electronic devices behind, as they are not permitted at the test site. For information on what can and cannot be brought to the exam site, see the CRISC exam candidate guide at `www.isaca.org/credentialing/crisc`.
- **Make a notification decision:** Decide whether you want your test results emailed to you. You will have the opportunity to consent to email notification of the exam results. If you are fully paid (zero balance on exam fee) and have agreed to the email notification, you should receive a one-time email approximately eight weeks from the date of the exam with your results.
- **Sleep:** Make sure you get a good night's sleep before the exam. Research suggests that you avoid caffeine at least four hours before bedtime, keep a notepad and pen next to the bed to capture late-night thoughts that might keep you awake, eliminate as much noise and light as possible, and keep your room at a suitable temperature for sleeping. In the morning, rise early so as not to rush and subject yourself to additional stress.

Day of the Exam

On the day of the exam, follow these tips:

- **Arrive early:** Check the Bulletin of Information and your confirmation letter for the exact time you are required to report to the test site. The confirmation letter or the candidate's guide explains that you must be at the test site *no later* than approximately 30 minutes *before* testing time. The examiner will begin reading the exam instructions at this time, and any latecomers will be disqualified from taking the test and will *not* receive a refund of fees.
- **Observe test center rules:** There may be rules about taking breaks. Rules will be discussed by the examiner, along with exam instructions. If you need anything at any time during the exam and are unsure of the rules, be sure to ask first. For information on conduct during the exam, see the ISACA CRISC candidate guide at `www.isaca.org/credentialing/crisc`.
- **Answer all exam questions:** Read the questions carefully, but do not overanalyze them. Remember to select the *best* solution. There may be several reasonable answers, but one is *better* than the others. If you aren't sure about an answer, you can mark the question and return to it later. After going through all the questions, you can return to the marked questions (and others) to read them and consider them more carefully. Above all, don't try to overanalyze questions, and do trust your instincts. Do not try to rush through the exam, as there is plenty of time to take as long as a few minutes on each question. However, at the same time, do watch the clock so that you don't find yourself going too slowly that you won't be able to answer every question thoughtfully.

- **Note your exam result:** When you have completed the exam, you should see your pass/fail result. Your results may not be in large, blinking text; you may need to read the fine print to get your preliminary results. If you passed, congratulations! If you did not pass, do observe any remarks about your status; you will be able to retake the exam—there is information about this on the ISACA website.

If You Did Not Pass

Don't lose heart if you did not pass your exam on the first attempt. Instead, remember that failure is a steppingstone to success. Thoughtfully take stock and determine your improvement areas. Go back to this book's practice exams and be honest with yourself regarding those areas where you need to learn more. Reread the chapters or sections where you need to know more. If you participated in a study group or training, contact your study group coach or class instructor if you feel you can get any advice from them on how to study up on the topics you need to master. Take at least several weeks to study those topics, refresh yourself on other topics, then give it another go. Success is granted to those who are persistent and determined.

After the Exam

A few weeks from the exam date, you will receive your exam results by email or postal mail. Each job practice area score will be noted in addition to the overall final score. All scores are scaled. Should you receive a passing score, you will also receive the application for certification.

Those unsuccessful in passing will also be notified. These individuals will want to closely look at the job practice area scores to determine areas for further study. They may retake the exam as many times as needed on future exam dates, as long as they have registered and paid the applicable fees. Regardless of pass or fail, exam results will not be disclosed via telephone, fax, or email (except for the consented email notification).



You are *not* permitted to display the CRISC moniker until you have completed certification. Passing the exam is not sufficient to use the CRISC designation anywhere, including in email, résumés, CVs, correspondence, or on social media.

Applying for CRISC Certification

To apply for certification, you must submit evidence of a passing score and related work experience. Keep in mind that you have five years to use this score on a CRISC certification application once you receive a passing score. After this time, you will need to retake the

exam. In addition, all work experience submitted must have been within ten years of your new certification application.

To complete the application process, you need to submit the following information:

- **CRISC application:** Note the exam ID number, as found in your exam results letter, list the information security management experience and any experience substitutions, and identify which CRISC job practice area (or areas) your experience pertains to.
- **Verification of Work Experience forms**: These must be completed and signed by your immediate supervisor or a person of higher rank in the organization to verify the work experience noted on the application. You must complete a set of Verification of Work Experience forms for each employer.

As with the exam, after you've successfully mailed the application, you must wait approximately eight weeks for processing. If your application is approved, you will receive an email notification, followed by a package in the mail containing your letter of certification, certificate, and a copy of the Continuing Professional Education Policy. You can then proudly display your certificate and use the "CRISC" designation on your résumé, email, social media profiles, and business cards.



You may use the CRISC moniker *only* after receiving your official certification letter from ISACA.

Retaining Your CRISC Certification

There is more to becoming a CRISC professional than merely passing an exam, submitting an application, and receiving a paper certificate. Becoming a CRISC professional is an ongoing and continuous lifestyle. Those with CRISC certification agree to abide by a code of ethics, meet ongoing education requirements, and pay annual certification maintenance fees. Let's take a closer look at the education requirements and explain the costs involved in retaining certification.

Continuing Education

The goal of continuing professional education requirements is to ensure that individuals maintain CRISC-related knowledge to better develop and manage security management programs. To maintain CRISC certification, individuals must complete 120 continuing education hours within three years, with a minimum of 20 hours per year. Each CPE hour is equal to 50 minutes of active participation in educational activities.

What Counts as a Valid CPE Credit?

For training and activities to be used for CPEs, they must involve technical or managerial training directly applicable to information security and information security management.

The following list of activities has been approved by the CRISC certification committee and can count toward your CPE requirements:

- ISACA professional education activities and meetings.
- Non-ISACA professional education activities and meetings.
- Self-study courses.
- Vendor sales or marketing presentations (ten-hour annual limit).
- Teaching, lecturing, or presenting on subjects related to job practice areas.
- Publication of articles and books related to the profession.
- Exam question development and review for any ISACA certification.
- Passing related professional examinations.
- Participation in ISACA boards or committees (20-hour annual limit per ISACA certification).
- Contributions to the information security management profession (ten-hour annual limit).
- Mentoring (ten-hour annual limit).

For more information on what is accepted as a valid CPE credit, see the Continuing Professional Education Policy (`www.isaca.org/credentialing/how-to-earn-cpe/#cpe-policy`).

Tracking and Submitting CPEs

Not only are you required to submit a CPE tracking form for the annual renewal process, but you should also keep detailed records for each activity. Records associated with each activity should have the following:

- Name of attendee
- Name of sponsoring organization
- Activity title
- Activity description
- Activity date
- Number of CPE hours awarded

It is in your best interest to track all CPE information in a single file or worksheet. ISACA has developed a tracking form for your use, which is available in the Continuing Professional Education Policy. To make it easier on yourself, consider keeping all related records, such as receipts, brochures, and certificates, in one place, like the example shown in Table 2. Documentation should be retained throughout the three-year certification period and for at least one additional year afterward. Evidence retention is essential, as you may someday be audited. If this happens, you will be required to submit all paperwork. So why not be prepared?

For new CRISCs, the annual and three-year certification periods begin on January 1 of the year following certification. You are not required to report CPE hours for the first partial year after your certification; however, the hours earned from the time of certification to December 31 can be utilized in the first certification reporting period the following year.

TABLE 2 Sample CPE records.

Name Chris Jacobs

Certification Number 67895787

Certification Period 1/1/2026 to 12/31/2028

Activity Title/Sponsor	Activity Description	Date	CPE Hours	Support Docs Included?
ISACA presentation/lunch	PCI compliance	2/11/2026	1 CPE	Yes (receipt)
ISACA presentation/lunch	Security in SDLC	3/21/2026	1 CPE	Yes (receipt)
Regional Conference, RIMS	Compliance, risk	1/12–14/2026	6 CPEs	Yes (CPE receipt)
BrightFly webinar	Governance, risk, & compliance	2/16/2026	3 CPEs	Yes (confirmation email)
ISACA board meeting	Chapter board meeting	4/8/2026	2 CPEs	Yes (meeting minutes)
Presented at ISSA meeting	Risk management presentation	6/21/2026	1 CPE	Yes (meeting notice)
Published an article in XYZ	Journal article on SOX ITGCs	4/12/2026	4 CPEs	Yes (article)
Vendor presentation	Learned about GRC tool capability	5/12/2026	2 CPEs	Yes
Employer-offered training	Change management course	3/25/2026	7 CPEs	Yes (certificate of course completion)

Therefore, should you get certified in January, you will have until the following January to accumulate CPEs. You will not have to report them until you report the totals for the following year, in October or November. This is known as the *renewal period*. During this time, you will receive an email directing you to the website to enter CPEs earned over the course of the year. Alternatively, the renewal will be mailed to you, and then CPEs can be recorded on the hard-copy invoice and sent with your maintenance fee payment. CPEs and maintenance fees must be received by January 15 to retain certification.

A notification of compliance from the certification department is sent after all information has been received and processed. Should ISACA have any questions about the information you have submitted, it will contact you directly.

CPE Maintenance Fees

To remain CRISC certified, you must pay CPE maintenance fees each year. These fees are (as of early 2026) $45 for members and $85 for nonmembers each year. These fees are in addition to ISACA membership and local chapter dues (neither of which is required to maintain your CRISC certification).

Revocation of Certification

A CRISC-certified individual may have their certification revoked for the following reasons:

- Failure to complete the minimum number of CPEs during the period.
- Failure to document and provide evidence of CPEs in an audit.
- Failure to submit payment for maintenance fees.
- Failure to comply with the Code of Professional Ethics can result in investigation and ultimately can lead to revocation of certification.

If you have received a revocation notice, you will need to contact the ISACA Certification Department at `certification@isaca.org` for more information.

Living the CRISC Lifestyle

Being a CRISC involves a lot more than passing the exam, participating in continuous learning, and paying the annual maintenance fees. There are numerous opportunities to get involved in local, national, and global activities and events that will help you grow professionally and meet other risk management professionals.

Find a Local Chapter

ISACA has over 200 local chapters in nearly 100 countries worldwide. Chances are, there is a chapter near you. I attended many ISACA chapter meetings and other events in Seattle when I lived there, where engaging speakers spoke on new topics, and where I met many like-minded security and audit professionals over the years.

Local chapters rely entirely on volunteers, and there is room for you to help in some way. Better chapters offer a variety of programs, events, study groups, and other activities that enrich participants professionally. For me, most of my ISACA experience happens in my local chapter.

Attend ISACA Events

ISACA hosts fantastic in-person conferences featuring world-class keynote speakers, expert presentations, vendor demonstrations and exhibits, a bookstore, and opportunities to meet other security, risk, and audit professionals. I find ISACA conferences enriching to the point

of being overwhelming. There are so many learning and networking opportunities that I find myself nearly exhausted at the end of an ISACA conference.

Join the Online Community

ISACA has an online community known as Engage, in which participants can discuss any topic related to security, risk, audit, privacy, and IT management. You can read and participate in online discussions, ask questions, help others with their questions, and make new professional connections. You can join Engage at `https://engage.isaca.org/`.

Pay It Forward Through Mentorship

If you are at the point in your career where you qualify for and have a reasonable prospect of passing the CRISC exam, chances are you have had a mentor or two earlier in your career, and maybe you have one now. As you grow in your professional stature, others will look to you as a potential mentor. Perhaps someone will come out and ask you to consider mentoring them.

The world needs more, and better, information security professionals and leaders. Mentoring is a great way to "pay it forward" by helping others get into the profession and grow professionally. You will also find that mentoring will enrich you as well.

Volunteer

As a nonprofit organization, ISACA relies on volunteers to enrich its programs and events. There are many ways to help, and one or more of these volunteer opportunities might be suitable for you:

- **Speaking at an ISACA event:** Whether you give a keynote address or a session on a specific topic, it's a mountaintop experience. You can share your knowledge and expertise on a particular topic with attendees, but you'll learn some things, too.
- **Serving as a chapter board member:** Local chapters don't run by themselves—they rely on volunteers who are working professionals who want to improve the lot of other professionals in the local community. Board members can serve in various ways, from financial management to membership to events.
- **Starting or helping a CRISC study group:** Whether as a part of a local chapter or at large, consider starting or helping a group of professionals who want to learn the details of the CRISC job practice. I am a proponent of study groups, because study group participants make the best students: they take the initiative to take on a big challenge to advance their careers.
- **Writing an article:** ISACA publishes online and paper-based publications with articles on a wide variety of subjects, including current developments in security, privacy, risk, and IT management from many perspectives. If you have specialized knowledge on some topic, other ISACA members can benefit from this knowledge if you write an article.

- **Participating in a credential working group:** ISACA works hard to ensure that its many certifications remain relevant and up to date. Experts around the world in many industries give their time to ensure that ISACA certifications remain the best in the world. ISACA conducts online and in-person working groups to update certification job practices, write certification exam questions, and publish updated study guides and practice exams. I contributed to the first working group in 2013, when ISACA initially developed the CRISC certification exam where I met many like-minded professionals, some of whom I am still in regular, meaningful contact with.
- **Participating in ISACA CommunITy Day:** ISACA organizes a global effort of local volunteering to make the world a better, safer place for everyone. Learn about the next CommunITy Day at `https://engage.isaca.org/communityday/`.
- **Writing certification exam questions**: ISACA needs experienced subject-matter experts willing to take the time to write new certification exam questions. ISACA has a rigorous, high-quality process for exam questions that includes training. Who knows—you could even be invited to an in-person workshop on writing exam items. You can find out more about how this works at `www.isaca.org/credentialing/write-an-exam-question`.

You can learn about these and many other volunteer opportunities at `https://engage.isaca.org/volunteeropportunities/about`.



Please take a minute to reflect on the quality and depth of ISACA and its many world-class certifications, publications, and events. These are all fueled by volunteers who made ISACA into what it is today. Only through your contributions of time and expertise will ISACA continue to excel for future security, risk, privacy, and IT professionals. And one last thing you can only experience on your own: volunteering not only helps others but enriches you as well. Will you consider leaving your mark and making ISACA better than you found it?

Continue to Grow Professionally

Continuous improvement is a mindset and a lifestyle that is built into IT service management and information security—it's even a formal requirement in ISO/IEC 27001! I suggest that you periodically take stock of your career status and aspirations, be honest with yourself, and determine what mountain you will climb next. If needed, find a mentor who can guide you and give you solid advice.

While this may not immediately make sense to you, know this: helping others, whether through any of the volunteer opportunities listed previously or in other ways, will enrich you personally and professionally. We're not talking about feathers in your cap or juicy items on your résumé, but rather the growth in character and wisdom that results from helping and serving others, notably when you initiated the helping and serving.

Professional growth means different things to different people. Whether it's a better job title, more money, a better (or bigger or smaller) employer, a different team, more responsibility, or more certifications, embarking on long-term career planning will pay dividends. Take control of your career and your career path—this is yours to own and shape as you will.

Additional Study Tools

This book comes with a number of additional study tools to help you prepare for the exam. They include the following.



Go to `www.wiley.com/go/Sybextestprep` to register and gain access to this interactive online learning environment and test bank with study tools.

Sybex Test Preparation Software

Sybex's test preparation software lets you prepare with electronic test versions of the review questions from each chapter, the practice exam, and the bonus exam that are included in this book. You can build and take tests on specific domains, by chapter, or cover the entire set of CRISC exam objectives using randomized tests.

Electronic Flashcards

Our electronic flashcards are designed to help you prepare for the exam. Over 100 flashcards will ensure that you know critical terms and concepts.

Glossary of Terms

Sybex provides a full glossary of terms in PDF format, allowing quick searches and easy reference to materials in this book.

Bonus Practice Exams

In addition to the practice questions for each chapter, this book includes two full 150-question practice exams. I recommend that you use them both to test your preparedness for the certification exam.



Like all exams, the Certified in Risk and Information Systems Control (CRISC) from ISACA is updated periodically and may eventually be retired or replaced. At some point after ISACA no longer offers this exam, the old editions of our books and online tools will be retired. If you have purchased this book after the exam was retired or are attempting to register in the Sybex online learning environment after the exam was retired, please know that we make no guarantees that this exam's online Sybex tools will be available once the exam is no longer available.

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. 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."

Assessment Test

Are you wondering how much you already know? Then go ahead and take this quick assessment test. Answers and explanations follow (no peeking) to see how you score.

1. Which activity best reflects effective organizational governance in IT risk management?
   - **A.** Selecting security tools based on industry trends
   - **B.** Aligning IT risk activities with business objectives
   - **C.** Performing penetration testing annually
   - **D.** Delegating all risk decisions to IT management
2. Which factor most directly influences the organization's risk appetite?
   - **A.** Number of identified vulnerabilities
   - **B.** Regulatory enforcement activity
   - **C.** Business strategy and objectives
   - **D.** Frequency of control testing
3. In threat modeling, what is the primary purpose of identifying threat agents?
   - **A.** Understanding sources and capabilities of threats
   - **B.** Assigning control ownership
   - **C.** Quantifying residual risk
   - **D.** Determining compliance gaps
4. What is the primary output of a business impact analysis (BIA)?
   - **A.** Risk appetite statement
   - **B.** Control testing results
   - **C.** Threat likelihood estimates
   - **D.** Impact and recovery requirements
5. Which situation most clearly indicates misalignment between risk appetite and risk tolerance?
   - **A.** Risks exceeding tolerance but approved through formal exceptions
   - **B.** Risks documented in the risk register
   - **C.** Controls mapped to multiple frameworks
   - **D.** KRIs reported quarterly
6. Which activity best validates that risk responses were implemented as approved?
   - **A.** Threat modeling
   - **B.** Risk reassessment
   - **C.** Control testing and monitoring
   - **D.** Policy review

7. Which reporting method best supports executive-level risk decision-making?
   - **A.** Detailed vulnerability scan reports
   - **B.** Heatmaps and dashboards
   - **C.** Technical configuration baselines
   - **D.** Incident response playbooks

8. Which principle ensures personnel understand their responsibilities for protecting information assets?
   - **A.** Enterprise architecture
   - **B.** Security awareness training
   - **C.** Data classification
   - **D.** Change management

9. Which scenario best represents a control deficiency?
   - **A.** A risk with high inherent impact
   - **B.** A policy not formally approved
   - **C.** A control that exists but fails to operate as intended
   - **D.** A threat actor targeting a system

10. What is the primary purpose of assigning risk ownership?
    - **A.** Ensuring audit independence
    - **B.** Establishing accountability for risk decisions
    - **C.** Centralizing control testing
    - **D.** Reducing inherent risk

11. Which metric is most appropriate for monitoring changes in risk exposure over time?
    - **A.** KPI
    - **B.** KCI
    - **C.** KRI
    - **D.** SLA

12. Which risk response option involves accepting risk without additional mitigation?
    - **A.** Avoid
    - **B.** Mitigate
    - **C.** Transfer
    - **D.** Accept

13. Which control type is intended to identify undesirable events after they occur?

- **A.** Preventive
- **B.** Detective
- **C.** Corrective
- **D.** Compensating

14. When evaluating residual risk, which factor is most critical?

- **A.** Number of implemented controls
- **B.** Control design intent
- **C.** Effectiveness of controls in operation
- **D.** Historical loss data only

15. Which IT discipline ensures systems support business requirements throughout their lifecycle?

- **A.** IT operations management
- **B.** SDLC
- **C.** Disaster recovery management
- **D.** Enterprise architecture

16. Which document is primarily used to record identified IT risks and their attributes?

- **A.** Control matrix
- **B.** Incident log
- **C.** Business impact analysis
- **D.** Risk register

17. Which risk response is MOST appropriate when a third party introduces systemic risk that cannot be fully mitigated internally?

- **A.** Accept
- **B.** Avoid
- **C.** Transfer
- **D.** Mitigate

18. Which condition most strongly justifies escalation of IT risk to enterprise risk management (ERM)?

- **A.** Risk exceeding business-defined tolerance thresholds
- **B.** High-likelihood, low-impact technical risk
- **C.** Lack of documented controls
- **D.** Frequent operational incidents

19. Which activity best supports management of emerging risk?

- **A.** Annual compliance audits
- **B.** Static risk registers
- **C.** Continuous environmental scanning
- **D.** Control self-assessments

20. What is the primary objective of issue, finding, and exception management?

- **A.** Eliminating all residual risk
- **B.** Ensuring deviations are tracked and resolved
- **C.** Replacing formal risk assessments
- **D.** Delegating accountability to auditors

Answers to Assessment Questions

1. B. Organizational governance requires alignment of IT risk management with business strategy, goals, and objectives. When they are aligned, IT risk decisions made by management are more likely to reflect the organization's risk appetite. Security tools should be selected based on identified risks and local IT standards. Annual penetration tests are essential, but not an indicator of business alignment. Risk decisions are made by business leaders, and rarely, if ever, by IT management.

2. C. Risk appetite is defined by leadership, based on strategic objectives and acceptable risk levels. Risk appetite is not influenced by the number of vulnerabilities or the frequency of control testing. Regulatory enforcement could influence risk appetite, but it is not the best answer.

3. A. Threat agents are identified to understand their motivations, who or what may exploit vulnerabilities, and how they do so. Identifying threat agents is unrelated to control ownership, residual risk, and compliance gaps.

4. D. A BIA identifies the most critical business processes and their dependencies, business impacts, and defines recovery priorities, such as recovery time objectives (RTOs) and recovery point objectives (RPOs). A BIA is unrelated to risk appetite, control testing results, and threat estimates.

5. A. Approving risks beyond tolerance without revisiting appetite indicates governance misalignment. Risks in the risk register, mapped controls, and KRI reporting are normal risk management activities and do not indicate misalignment.

6. C. Control testing and monitoring confirm that approved risk responses are operating effectively. While essential, threat modeling and policy reviews are unrelated. Risk reassessments may serve as validation, but control testing and monitoring are the preferred methods.

7. B. Aggregated visual reporting supports timely, risk-based decisions by senior management. Vulnerability scan reports, configuration baselines, and incident response playbooks, while essential, are unrelated to executive-level risk decision-making.

8. B. Security awareness training promotes a risk-aware culture by educating personnel on information security responsibilities. Enterprise architecture, data classification, and change management are all essential activities that staff can learn from, but are not as effective as security awareness training.

9. C. Control deficiencies occur when controls are ineffective or not operating as designed. The other choices represent risk but are not necessarily control deficiencies, as there may be no controls in these areas.

10. B. Risk ownership establishes accountability for managing and responding to identified risks. Risk ownership is unrelated to audit independence and control testing and does not influence inherent risk.

11. C. Key Risk Indicators (KRIs) are designed to signal changes in risk levels. Key Performance Indicators (KPIs) and Key Control Indicators (KCIs) are essential in risk management but are not well-suited to monitoring changes in risk exposure. Service Level Agreements (SLAs) are not used to monitor changes in risk; they define service response levels.

12. D. Risk acceptance occurs when management determines that the risk is within its appetite and does not require further action. Risk avoidance involves discontinuing the activity that introduced the risk. Risk mitigation involves changes to processes or systems to reduce risk. Risk transfer involves the process of transferring risk to an external party.

13. B. Detective controls identify and report events that have already occurred. Preventive controls are designed to keep unwanted events or conditions from occurring. Corrective and compensating controls are used to correct a condition and compensate for the absence or weakness of another control, respectively.

14. C. Residual risk depends on how effectively controls reduce inherent risk in practice. The number of controls and their intent are unrelated. Historic data loss may inform future decisions, but it is not as critical as the effectiveness of controls in operation.

15. B. The systems (or software) development lifecycle (SDLC) governs system design, development, testing, implementation, and maintenance. IT operations management is essential, but it does not support business requirements throughout a system's lifecycle. Disaster recovery management and enterprise architecture are generally essential, but they are unrelated.

16. D. The risk register is the authoritative record for documenting risks, including likelihood, impact, ownership, and status. A control matrix is like a heatmap and is used to depict risks visually. An incident log and a business impact analysis (BIA) contain helpful information, but are not used to record IT risks.

17. C. Risk transfer, such as through contractual terms or insurance, is appropriate for third-party risks that cannot be fully mitigated internally. Accepting risk leaves the systemic risk in place. Avoiding risk is often impractical, as the organization would have to discontinue the activity(ies) that introduced the risk in the first place.

18. A. Risks exceeding tolerance thresholds require enterprise-level visibility and decision-making. The other risks shown are less likely to result in an escalation to ERM.

19. C. Emerging risks require ongoing monitoring of internal and external environments. Annual compliance audits, static risk registers, and control self-assessments are not designed to identify emerging risks.

20. B. These processes ensure identified issues and exceptions are documented, tracked, and remediated. Issue, finding, and exception management are essential activities in a risk management program for addressing these matters formally. They do little, if anything, to reduce or eliminate residual risk, and they do not replace formal risk assessments. Risk accountability is rarely, if ever, delegated to auditors.



Chapter 1

Governance

This chapter covers CRISC Domain 1, "Governance," including the Organizational Governance and Risk Governance subdomains. The job practice for this domain is described by ISACA as follows:

The governance domain interrogates your knowledge of information about an organization's business and IT environments, organizational strategy, goals and objectives, and examines potential or realized impacts of IT risk to the organization's business objectives and operations, including Enterprise Risk Management and Risk Management Framework.

The job practice consists of these elements:

✔ A—ORGANIZATIONAL GOVERNANCE

1. *Organizational Strategy, Goals, and Objectives*
2. *Organizational Structure, Roles, and Responsibilities*
3. *Organizational Culture*
4. *Policies and Standards*
5. *Business Processes*
6. *Organizational Assets*

✔ B—RISK GOVERNANCE

1. *Enterprise Risk Management and Risk Management Framework*
2. *Three Lines of Defense*
3. *Risk Profile*
4. *Risk Appetite and Risk Tolerance*
5. *Legal, Regulatory, and Contractual Requirements*
6. *Professional Ethics of Risk Management*

The CRISC Task Statements relevant to this domain focus on governance and how it applies in this context to the organization, particularly executive management.

The CRISC Task Statements covered in this chapter are:

1. ***Collect and review existing information regarding the organization's business and IT environments.***
2. ***Identify potential or realized impacts of IT risk to the organization's business objectives and operations.***
7. ***Facilitate the identification of risk appetite and risk tolerance by key stakeholders.***
22. ***Evaluate alignment of business practices with risk management and information security frameworks and standards.***

The topics in this chapter represent 26% of the CRISC examination. This chapter explores the governance aspects of risk for an organization, as well as the relationships between enterprise risk and IT risk. The chapter also examines risk management governance and its implementation within the organizational structure and culture. Governance also encompasses overall laws, regulations, and risk frameworks, which are discussed in detail. During these governance discussions, both external governance and internal governance, including their policies and standards, will be explained.

Internal governance also encompasses management oversight, including business process reviews, risk appetite, and tolerance, as well as contractual obligations related to risk management. Other topics related to the Task Statements touched on here and examined in depth in other chapters include elements of risk, such as assets and their value, threats, vulnerabilities, likelihood, and impact.

Although many organizations will "do the right thing" in terms of due diligence and care, upholding their legal and ethical responsibilities, many other organizations don't do it to the same standards. This is why governance is crucial to an organization. Governance ties all these elements together. Governance sets standards—whether they are legal, ethical, or otherwise—and attempts to "standardize" what organizations do in terms of their obligations to stakeholders and society in general. As you'll see in this chapter, governance drivers originate from various sources, both within and outside the organization. It is a synthesis of enforced requirements and organizational leadership culture. Additionally, governance is essential to ensure that organizations handle systems and data with care and fulfill their due diligence and due care obligations when managing data protection and risk.

This chapter focuses on organizational governance and its application to all the various elements of an information security management system, including mission, strategy, goals and objectives, roles and responsibilities, culture, risk, and ethics.

Organizational Governance

Governance is the glue that holds all the different elements of an organization together, including its mission, strategy, goals, objectives, and operations. Governance establishes the policies and requirements an organization must meet and consists of both internal and external governance. External inputs come in the form of laws, regulations, professional and

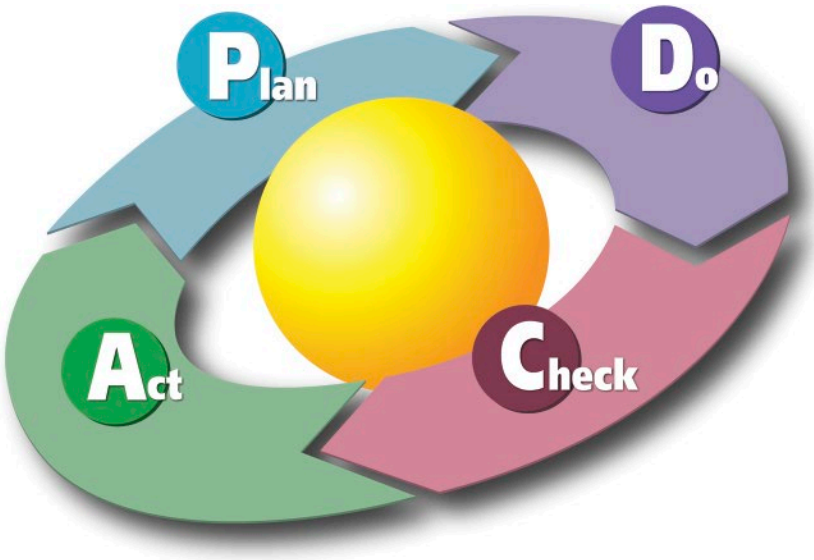
industry standards, and other sources of requirements imposed on the organization from the outside. Internal inputs typically support external inputs through policies, procedures, and processes. For example, if a law imposes requirements to protect certain sensitive data using specific controls or to a certain standard, then policies and procedures further support those requirements by formalizing and codifying them within the organization. Policies and procedures can also be independent of external governance and simply reflect the organization's culture, needs, and values, usually established by its executive management. In any event, governance is the controlling factor for the organization; it keeps the organization on the right track and ensures that it meets its compulsory requirements, such as adherence to laws, as well as fulfilling its due diligence and due care responsibilities.

In addition to imposing requirements on the organization, governance also refers to the structure by which the organization is led and regulated. Again, this can come from external or internal sources. External drivers for governance may include an external board of directors or regulatory agencies, which ensure that the organization is led with a focus on responsibility and accountability. Internal governance drivers originate from internal business leaders and reflect business drivers unrelated to externalities. In reality, good governance typically demonstrates a balance between internal and external business drivers.

Industry IT governance frameworks serve as models for organizations building or improving their governance structures. COSO ERM 2017, NIST RMF, ISO 31000, and COBIT 2019 are governance framework models that outline standard approaches for aligning IT strategies with business strategies to maintain shareholder value. Governance is a *lifecycle* process, modeled after the *Deming Cycle*, which is depicted in Figure 1.1.

The organizational infrastructure framework, encompassing strategy, goals, objectives, mission statements, and other elements, supports governance, as discussed further in this chapter.

FIGURE 1.1 The Deming Cycle depicts a typical lifecycle process.



Source: Image from Karn Bulsuk (www.bulsuk.com).

Organizational Strategy, Goals, and Objectives

Every organization operates with a *mission* that defines its purpose and long-term direction. The organization is in business for a particular purpose, not merely because people want to come to work every day to perform tasks and socialize. Missions, goals, and objectives directly relate to why the organization is in business in the first place, whether that is to develop and market a product or provide a service. Organizational senior management defines the business's mission, goals, and objectives, typically on a strategic or long-term level. Senior management also determines the levels of risk tolerance, appetite, and capacity (discussed later in this chapter), based on factors such as the market space, operational environment, economy, and governmental regulations. These risk levels directly align with and support the business mission, balancing business opportunities that can generate revenue and move the business forward with potential adverse events that may cause the business to fail or at least have a detrimental impact on the organization.

One of the critical characteristics of governance structures is their *business alignment*, which refers to the degree to which those structures support the organization's mission, goals, policies, and culture. Processes and policies that are business-aligned will work with better harmony in the organization than those that are not.



Risk appetite and tolerance are directly related to the business mission, although different business pursuits may have varying levels of each. Senior management sets those levels based on the potential rewards from risky opportunities and the amount of loss the organization could endure if those rewards don't materialize.

Organizational Structure, Roles, and Responsibilities

How the business is organized can help drive how it deals with risk in several ways. Most businesses are organized from a functional perspective; in other words, departments and other hierarchical structures are established to handle specific functions that contribute to the business's goals and objectives. Also, organizational structure determines how authority and responsibility are distributed.

An organization's board of directors provides ultimate oversight and often establishes specialized committees such as audit, risk, technology, or ESG (environmental, social, and governance). Executive management implements governance through an enterprise risk and control structure. The Chief Executive Officer (CEO) aligns strategy with risk appetite, the Chief Information Officer (CIO) and Chief Information Security Officer (CISO) manage technology and security risk, the Chief Risk Officer (CRO) oversees the enterprise risk framework, and the Data Protection Officer (DPO) manages data and privacy compliance. In technology-driven organizations, a Chief AI Officer (CAIO) may oversee responsible artificial intelligence (AI) practices.

For example, in a production-driven business, there may be a manufacturing or production department, an engineering department, a research and development department,

and an assembly line. There will likely be additional departments that cover support functions, such as marketing, accounting and finance, public relations, and so on. A hospital, on the other hand, will be organized according to its specific functions, such as the emergency department, surgery, neurology, radiology, and so on. Businesses in other markets or areas will be organized differently as well. In any case, the organization of the business is structured as its mission and business purposes dictate. Certain functions can be found in any business, such as information technology, information security, and even legal compliance. These functional areas may have the primary function of dealing with risk. Still, it is also essential to consider that all organizational structures, from lower-level work sections to higher-level departments and divisions, have responsibilities regarding risk management.

The organization must examine its structure and determine how each unit will manage risk at its own level, recognizing that risk management should be uniform throughout the entire organization and explicit at all levels. Business unit leaders identify and manage risks within their domains, while risk and compliance professionals provide guidance and oversight. For example, risks that the accounting department incurs are only a subset of the risks at higher organizational levels and are incorporated into their risk management processes. Each lower-level unit in an organizational hierarchy has risks that are part of the next higher level's risk considerations. While individual units may be responsible for only a small piece of the overall organizational risk, their parent units also bear responsibility for managing that risk, as well as the risk of other subordinate units. Another concept related to organizational structures is that all parts of the organization bear the risk incurred by any part of the organization; there is almost no risk that only affects a small part of the business. Risk ripples across the entire organization in some way.

Each unit, whether it is located in the lower levels of the business hierarchy or at the highest levels, should take steps to identify, evaluate, and assess risks at its respective level. Risks can be categorized into tactical, operational, and strategic levels. *Tactical risks* are those short-term risks that are encountered by smaller-level production sections—those that carry out the day-to-day work of the organization. *Operational risks* are medium-term risks that can span multiple work units and relate to how the business conducts its functions, as well as how the various work units interact. *Strategic risk* is long-term risk borne at the higher levels of the organization, including senior management. It involves risks incurred by leading the business toward opportunities and away from decisions that exceed the organization's risk appetite and tolerance. Respectively, these three types of risks also correspond to short-term, medium-term, and long-term risks.

Regardless of the level of risk incurred within the organization, there should be an enterprise risk management strategy and program in place to deal with the lower-tier, middle-tier, and higher-tier risks as well as ensure that the risk management strategy handles risks consistently and uniformly in alignment with the organization's risk tolerance and appetite. Governance at higher levels of the organization influences risk appetite and tolerance, shaping the organization's risk management strategy throughout the organization. The organizational structure must support that governance, as well as clearly define lines of authority and responsibility in terms of risk leadership and management.

Selecting Stakeholders for the Cyber Risk Governance Committee

A software-as-a-service (SaaS) organization recently hired a new CISO to make strategic improvements to cybersecurity risk management. Experienced in risk management and risk governance, the CISO assessed the current state of cybersecurity and risk management.

The CISO interviewed senior executives and department heads to understand their opinions and concerns about cyber risk management in the organization. The CISO needed to know how the organization works, as well as the attitudes of those interviewed.

The CISO identified key *stakeholders*, including the CIO, heads of HR, legal, product development, customer support, and finance, and explained to each of them how cyber risk management could be implemented within the organization without disrupting critical business processes.

Part of the CISO's objective in these meetings was to learn from these leaders, and to build trust with them. Without trust, new processes like cyber risk management are unlikely to succeed.

Organizational Culture

Organizational *culture* is the term that describes how people treat each other, how they make decisions, and how they respond to risk. Many organizations establish a set of values that define the norms of behavior. Terms such as *respect*, *collaboration*, and *teamwork* are often reflected in these values. Some organizations publish formal *value statements* and display them in lobbies, offices, and conference rooms. Many also publish a code of conduct that defines the boundaries of ethical behavior.

The way that an organization's leaders treat each other and the rest of the organization sets an example for behavioral norms. Often, these norms reflect those formal values, but sometimes they may differ. One could say that an organization's stated culture and its actual culture may vary a little, or a lot. The degree of alignment itself reflects an organization's culture.

An organization also has a *risk culture*, which is essentially how the organization as an entity feels about and deals with risk. This culture is developed from several sources. First, it can come from the organization's leadership, based on their business and management philosophies, attitudes, education, and experience. It can also come from the organization's governance. Remember that governance is essentially the rules and regulations imposed on the organization by either external entities (in the form of laws, for example) or internally by the organization itself. In any case, the culture of the organization really defines how the organization feels about risk and how it treats risk over time. You will read later about how these two concepts—risk tolerance and risk appetite—support the organizational risk culture.

Policies and Standards

Policies are the instruments through which governance intent becomes operational reality. They formalize management's directives regarding acceptable behavior, security requirements, and control expectations.

The most common policy hierarchy typically includes:

- ***Policy*:** Broad management statements of intent
- ***Standard*:** Mandatory requirements or specifications that support policy
- ***Process*:** A grouping of procedures that support a business function
- ***Procedure*:** Step-by-step instructions for performing tasks
- ***Guideline*:** Recommended practice for implementing a policy

Controls exist outside the policy hierarchy but are mapped to policies or derived from them. Different types of controls work together to protect an organization, many of which have a technical focus. However, while the more technical approaches are often the stars of the cybersecurity show, the importance of good old-fashioned policy cannot be overstated. Controls are discussed in detail in Chapter 3, "Risk Response and Reporting."

Without a policy, a security program often falls short due to a lack of governance, management involvement, and accountability. This objective covers different types of policies and how they work together to form the underlying foundation of human and process control, ensuring a balanced and effective security program.

To provide effective security, security policy and procedure creation must begin at the top of an organization, with senior management that understands, supports, and holds personnel accountable for adhering to the written policies. Remember that organizational policies must articulate governance, reflecting the requirements found in laws and regulations, as well as the management philosophies and tolerances of the organization's leadership. These policies and procedures must then flow throughout the organization to ensure that security is valued and functional at every level. Understanding organizational security must begin with an understanding of the fundamental laws, regulations, and legal jurisdiction that policies must be informed by to protect not only the organization and its assets but also its employees and customers.

Management creates policies based on its leadership philosophies, culture, industry standards and requirements, and, most importantly, external governance inputs. Policies articulate with all these items and represent management's directives in upholding these requirements. Policies specific to risk governance include a body of documents that may begin with the risk management strategy, which outlines the organization's overall strategy, methodology, and long-term objectives for managing risk. Policies support the strategy, as well as governance requirements, such as those that may be found in laws or regulations that dictate that certain risk management activities must take place within the organization. An example of a risk-related policy is a risk management or risk assessment policy, which outlines the organization's requirements for risk management, assessment, analysis, treatment, and monitoring. This policy may be brief or detailed, but will usually state roles and responsibilities, the risk management methodology adopted by the organization, and the requirement to implement detailed risk management procedures.

Policy governance follows a defined lifecycle: creation, approval, publication, review, and retirement. Policies must be endorsed by senior management and communicated across all levels. They should be reviewed periodically (annually is a good start), particularly when regulations, technologies, or business conditions change.

Modern organizations use automated policy management tools that track control status, regulatory mapping, and version history. Continuous compliance programs integrate policy enforcement with enterprise *integrated risk management (IRM)* (formerly known as governance, risk, and compliance [GRC]) systems.

Policies should align with external frameworks such as ISO/IEC (International Organization for Standardization and the International Electrotechnical Commission) 27001:2022, ISO/IEC 42001 (Artificial Intelligence Management System), National Institute of Standards and Technology (NIST) Cybersecurity Framework 2.0, and/or other relevant frameworks. Mapping internal policies to these standards improves consistency and audit readiness.

Business Processes

Business processes are structured activities that fulfill the organization's mission. Governance ensures that these processes operate efficiently, meet objectives, and remain within established risk boundaries.

Each process has an identified owner responsible for outcomes and associated risks. Process governance requires documentation of roles, dependencies, and control activities.

For example, a shoe manufacturer is in the business of making shoes; therefore, its mission is to make good-quality shoes. Most businesses have a mission statement that describes their mission. While the mission is the overall reason for existence, business processes are the activities that carry out that mission. A shoe manufacturer's business processes could be as high-level as manufacturing, sales, and marketing, but even those higher-level processes are broken down into smaller activities, such as sewing cloth and leather, developing types and styles of shoes, and selling them to retailers. These are all processes that support the business.

Each of these business processes also incurs some level of business risk. There could be risks in the manufacturing process due to the types of materials or a specific manner of attaching the shoe soles to the rest of the shoe. A faulty process could mean that the shoes are of lesser quality than other competitors or that the styles don't meet the demands of consumers, leading to fewer sales or more consumer returns. Regardless of whether they are higher-level processes or more detailed activities and tasks, each is supported by one or more systems or sets of data. These systems also incur risk, but of a different nature. The risk incurred by the systems includes the ability to effectively use those systems and data to support various business processes. Therefore, IT (or even cybersecurity) risk directly informs business process risk, as it can affect the organization's ability to carry out its business processes and, in turn, its overall mission.

Business processes should be "owned" by different managers within the organization, which means that business process owners are responsible for both the day-to-day and the long-term operations and success of the process. This also means that they are ordinarily

the primary owners of the business risk associated with the process. Of course, higher-level management also bears ownership and responsibility for both the process and risk; however, these are typically the people in charge of the business process.

Automation technologies, including robotic process automation and AI-assisted decision systems, can increase efficiency but introduce governance challenges related to control validation and transparency. Governance frameworks ensure that automation remains aligned with policy, ethics, and compliance expectations.

Better organizations will take a formal approach regarding the development and management of business processes. Documents that describe processes, roles and responsibilities, and key assets may be formally managed through business process change management and periodic review, and kept in official repositories to ensure that no unauthorized changes occur. Process *maturity* and capability can be evaluated using models such as the COBIT Process Assessment Model. Process maturity is the degree of formality and optimization in place for a business process. Continuous improvement involves periodic review, automation, and optimization. COBIT is discussed in detail in Chapter 4.

Measuring business processes in various ways supports governance effectiveness:

- Key Performance Indicators (KPIs) track output and efficiency.
- Key Risk Indicators (KRIs) measure exposure levels.
- Key Control Indicators (KCIs) assess whether controls perform as intended.

Organizations with higher maturity will develop metrics and KPIs associated with each business process, allowing management to understand the quantity and quality of business process output. This understanding can be used to identify tactical and strategic improvements that enhance the effectiveness and efficiency of processes. Risk practitioners must work together with business process owners to develop these KPIs, as well as KRIs and KCIs.



KPIs, as well as KRIs and KCIs, are covered in more detail in Chapter 3, which discusses risk response and reporting as part of CRISC Domain 3.

Organizational Assets

Company *assets* can include physical items, such as computer and networking equipment, office machines (for example, scanners and copiers), work facilities, and information processing centers, as well as less tangible assets, such as software, business records, configuration data, license and activation keys, login credentials, and data. Asset governance provides the foundation for effective risk management and control of information systems. Asset identification involves identifying both types of assets and determining their value. Asset values must be established beyond the mere capital costs; an accurate asset valuation should consider several factors. For example, a consideration should be the cost of repairing or recovering the asset versus simply replacing it outright. Often, repairing the asset may be less expensive in the short run, but the cost of the different components required to conduct a repair should be

considered. Also, it's important to remember that this might only be a temporary solution—one that could come back to haunt you (and your budget) in the long run.

Asset management encompasses a range of activities designed to oversee the inventory, classification, utilization, and disposition of assets. Asset management is a foundational activity, without which several other activities could not be effectively done, including vulnerability management, device hardening, incident management, data security (these terms are discussed in Chapter 4), and some aspects of financial management.

In information security, asset management is critical to the success of vulnerability and attack surface management, which is discussed in Chapter 4. If assets are not known to exist, they may be excluded from processes used to identify and remediate vulnerabilities. Similarly, it will be difficult to harden assets if their existence is not known. Furthermore, if an unknown asset is attacked, the organization may have no way of directly realizing this in a timely manner. If an attacker compromises an unknown device, the attack may not be known until the attacker pivots and selects additional assets to compromise. This time lag could prove crucial to the impact of the incident.

The criticality of asset management warrants governance emphasis on these activities.



Senior executives in many organizations still do not fully understand the importance of maintaining a complete and accurate inventory of assets in the context of cybersecurity. In other contexts, incomplete asset inventories pose quality and financial efficiency issues. However, in the context of cybersecurity, failing to maintain an accurate inventory can permit existential threats to manifest and potentially destroy the business.

Asset Identification

A security management program's primary objective (whether formally stated or not) is the protection of the organization's assets. These assets may be tangible or intangible, physical, logical, or virtual. Here are some examples of assets:

- **Buildings and property:** These assets include real estate, structures, and other improvements.
- **Equipment:** This can include machinery, vehicles, and office equipment such as copiers, printers, and scanners.
- **IT equipment:** This includes computers, printers, scanners, tape libraries (the devices that create backup tapes, not the tapes themselves), storage systems, network devices, and phone systems.
- **Virtual assets:** In addition to the tangible IT equipment cited, virtual assets include virtual machines and the software running on them.
- **SaaS applications:** While owned and operated by external vendors, they perform functions critical to the organization. *Shared responsibility models* and contractual terms define vendor and customer roles and vary by service provider.

- **Supplies and materials:** These can include office supplies as well as materials used in manufacturing.
- **Records:** These include business records, such as contracts, video surveillance tapes, visitor logs, and more.
- **Information:** This includes data in software applications, documents, email messages, and files of all kinds on workstations and servers.
- **Intellectual property:** This includes an organization's designs, architectures, patents, software source code, processes, and procedures.
- **Personnel:** In a real sense, personnel *are* the organization. Without its staff, the organization cannot perform or sustain its processes.
- **Reputation:** One of the intangible characteristics of an organization, reputation refers to the individual and collective opinion about an organization held by its customers, competitors, shareholders, and the broader community.
- **Brand equity:** Similar to reputation, this refers to the perceived or actual market value of a brand, product, or service produced by the organization.

Sources of Asset Data

An organization that is building or improving its security management program may need to build its asset inventory from scratch. Management will need to determine where this initial asset data will come from. Some sources include the following:

- **Financial system asset inventory:** An organization that keeps all of its assets on the books will have a wealth of asset inventory information. However, it may not be entirely useful: asset lists often do not include the location or purpose of the asset and whether it is still in use. Correlating a financial asset inventory with assets in actual use may require more effort than other methods for creating the initial asset list. However, for organizations that have a relatively small number of highly valued assets (for instance, an ore crusher in a gold mine or a mainframe computer), knowing the precise financial value of an asset is highly useful because the actual depreciated value of the asset is used in the risk analysis phase of risk management. Knowing the depreciated value of other assets is also useful, as this will figure into the risk treatment choices that will be identified later.



Financial records that indicate the value of an asset may not include the value of information stored on (or processed by) the asset.

- **Interviews:** Discussions with key personnel for the purpose of identifying assets are usually the best approach. However, to be effective, several people usually need to be interviewed to be sure to include all relevant assets.
- **IT systems portfolio:** A well-managed IT organization will have formal documents and records for its major applications. While this information may not encompass every IT

asset in the organization, it can provide information on the assets supporting individual applications or geographic locations.

- **Online data:** An organization with a large number of IT assets (systems, network devices, and so on) can sometimes utilize the capability of online data to identify those assets. An organization with cloud-based assets can utilize the asset management portion of the cloud services dashboard to determine the number and type of assets in use. Additionally, a systems or network management system often includes a list of managed assets, which can serve as a good starting point when creating the initial asset list.
- **Security scans:** An organization that has security scanning tools can use them to identify network assets. This technique will identify both authorized and unauthorized assets.
- **Asset management system:** Larger organizations may find it more cost-effective to use an asset management application dedicated to this purpose, rather than relying on lists of assets from other sources.

None of these sources should be considered accurate or complete. Instead, as a formal asset inventory is being assembled, the security manager should continue to explore other sources of assets.

Collecting and Organizing Asset Data

It is rarely possible to take (or create) a list of assets from a single source. Instead, more than one source of information is often needed to ensure that the risk management program has identified at least the critical, in-scope assets it needs to worry about.

When asset inventory is not prioritized, other risks emerge. From a consequence-based perspective, unidentified or unmanaged assets represent the highest level of criticality, since their existence outside of formal control boundaries, business processes, and tooling creates unknown exposure and unmitigated risk. While an unidentified asset may not play a critical role in business process support, its neglect may substantially increase risk—risk that may not be specifically identified, and yet it exists. In cybersecurity, this is informally known as *dark risk* or shadow risk.



As part of IT governance, management needs to determine who is responsible for maintaining asset inventories.

It is usually helpful to organize or classify assets, which helps identify assets into smaller chunks that can be analyzed more effectively. There is no single way to organize assets, but here are a few ideas:

- **Geography:** A widely dispersed organization may want to classify its assets according to their location, aiding risk managers during the risk analysis phase, since many risks are geographic-centric, particularly natural hazards.
- **Service provider:** An organization utilizing one or more infrastructure-as-a-service (IaaS) providers can group its assets by service provider.

- **Business process:** Because some organizations rank the criticality of their individual business processes, it can be useful to group assets according to the business processes they support. This classification helps the risk analysis and risk treatment phases because assets supporting individual processes can be associated with business criticality and treated appropriately.
- **Organizational unit:** In larger organizations, it may be easier to classify assets according to the organizational unit they support.
- **Sensitivity:** Usually ascribed to information, sensitivity relates to the nature and content of that information and its relative confidentiality. Sensitivity usually applies in two ways: to an individual, where the information is considered personal or private, and to an organization, where the information may be considered a trade secret. Sometimes sensitivity is somewhat subjective and arbitrary, but often it is defined in laws and regulations.
- **Criticality:** Whereas some assets may be classified according to their sensitivity, others may also be classified by their criticality (or both), which is their importance to the organization in terms of performing its mission. Not all assets are sensitive, but they may still be highly necessary for the organization to function. Conversely, not all assets are considered critical, but still could be classified as sensitive in terms of confidentiality. Note that assets could be considered either sensitive or critical, or both.
- **Regulation:** For organizations required to follow government and other legal obligations regarding the processing and protection of information, it will be helpful to include data points that indicate whether specific assets are considered in scope for specific regulations. This is important because some regulations specify how assets should be protected, so it's helpful to be aware of this during risk analysis and risk treatment.

There is no need to choose which of these methods will be used to classify assets. Instead, an analyst should collect several points of metadata about each asset (including location, process supported, and organizational unit supported). This will enable the security manager to sort and filter the list of assets in various ways to better understand which assets are in a given location or which ones support a particular process or part of the business.

Organizations use configuration management databases (CMDBs), discovery tools, and GRC-integrated asset registers to maintain visibility and accountability. Automated synchronization among asset, vulnerability, and compliance systems strengthens overall governance.

Risk Governance

Risk governance has several different aspects. As mentioned earlier, governance outlines the overall requirements that the organization must adhere to, which may include legal, regulatory, ethical, safety, and professional standards. Risk governance, then, describes the requirements that an organization must adhere to in managing both business and IT risk. Laws, regulations, and other external inputs can dictate portions of risk governance, as can internal governance, which comes in the form of policy. Risk governance is also established

by executive management through the formation of risk appetite and tolerance, risk strategy, and through various policies that support risk management. Risk governance encompasses adopting a risk assessment and analysis methodology, risk treatment and response options, risk monitoring processes, and so on. Table 1.1 lists the various facets of risk governance and their coverage in this book.

TABLE 1.1 The principal components of risk governance.

Topic	Description	Coverage
Risk culture and leadership	Norms and behaviors	Chapter 1
Enterprise risk management	Comprehensive risk on all topics, including cybersecurity	Chapter 1
Risk management frameworks	Industry standard risk management methodologies	Chapter 1
Three lines of defense	An industry standard assurance and accountability model	Chapter 1
Risk appetite and risk tolerance	Thresholds that guide risk treatment decisions	Chapter 1
Risk assessment and risk analysis	Industry standard risk assessment and risk analysis methodologies	Chapter 2
Risk response and risk treatment	Decisions and follow-up on the disposition of identified risks	Chapter 3
Control design and selection	The design, selection, and implementation of controls to reduce identified risks	Chapter 3
Risk monitoring	Continuous and periodic risk assessments	Chapter 3
Risk metrics and reporting	The development of metrics and reporting for senior management	Chapter 3
Continuous improvement	Periodic updates to risk governance processes	Chapter 3
Security policies, standards, and guidelines	Cybersecurity-related corporate policy defining expected behavior and characteristics	Chapter 5

ERM and Risk Management Frameworks

Enterprise risk management (ERM) is the practice of identifying and managing strategic risk in an organization. Topics of ERM include macroeconomics, market risk, regulations, workforce, information technology, and cybersecurity. Individual risks in ERM will cite various uncertainties and scenarios that, if realized, may impede the organization's long-term viability.

Standards and frameworks help provide a standardized, industry-accepted approach toward managing risk. While they're certainly not perfect, they provide a baseline for an organization to recognize the key steps toward understanding and categorizing the risks to its business, implementing the appropriate security controls to lower risk to an acceptable level, and continually monitoring the residual risk and reporting the status to management. Each standard and framework has its own pluses and minuses, and this section discusses two of the more commonly used: the ISACA Risk IT Framework and the National Institute of Standards and Technology (NIST) Risk Management Framework. Keep in mind that this discussion is intended to introduce you to the frameworks.

ISACA Risk IT Framework

Initially published in 2009, the ISACA *Risk IT Framework* is a widely recognized risk management framework you should be aware of. ISACA created a framework that merges the traditional IT models with a more risk-focused mindset. Within the ISACA model, the activities considered key are grouped into processes, which in turn are grouped into three domains. You will see many of the same concepts that you'll find within this book (and even this chapter), so it's definitely a worthwhile model to review. The domains are as follows:

- **Risk Governance:** Discusses risk appetite, risk tolerance, and communicating risk to leadership.
- **Risk Evaluation:** Includes an understanding of the impacts to the business and the risk scenarios, developed through the inclusion of internal and external environmental factors, risk management capability, IT capability, and IT-related business capability.
- **Risk Response:** Covers the risk response definition, risk response prioritization, and key risk indicators (KRIs).

The framework outlines risk response options, including risk avoidance, risk mitigation, risk sharing and transfer, and risk acceptance.

NIST Risk Management Framework

NIST Special Publication 800-37, Revision 2, "Guide for Applying the Risk Management Framework to Federal Information Systems," along with other supporting publications from NIST (part of the U.S. Department of Commerce), make up the defining volumes of the *NIST Risk Management Framework (RMF)*. The RMF is composed of seven steps:

1. Prepare.
2. Categorize information systems.
3. Select security controls.

4. Implement security controls.
5. Assess security controls.
6. Authorize information systems.
7. Monitor security controls.

Security categorization requires the information systems and associated information to be categorized based on an analysis of the impact to the organization of a loss of confidentiality, integrity, and availability for the system. When selecting a set of security controls, you must consider both the security categorization and any previous risk assessments. Following this step, you implement the previously selected security controls and then assess their effectiveness and the overall success of the controls in reducing the risk to the accepted level. Finally, security controls must be continually monitored to ensure they continue to operate as intended. You also need to consider any changes that come into effect across the organization (think location, mission, security classification, or other underlying factors).

You will notice many additional concepts within NIST Special Publication 800-37 that are germane to the CRISC exam; for example, the early steps of the RMF discuss categorizing and selecting appropriate security controls (discussed later in the book), while other steps correspond to the implementation, assessment, and continuous monitoring of controls (covered throughout this book). For these reasons alone, it's a great idea to read this document to understand how the steps work together and how they are similar to (and sometimes different than) other frameworks.

COSO Enterprise Risk Management

The *COSO Enterprise Risk Management (ERM) Framework* provides a structured approach for identifying, assessing, and managing risks in alignment with organizational strategy. It expands on COSO's internal control model to emphasize governance, culture, risk appetite, and performance integration. The framework consists of five interrelated components—governance and culture, strategy and objective-setting, performance, review and revision, and information, communication, and reporting—supported by 20 principles. COSO ERM promotes proactive risk management, enhancing decision-making and organizational value. The full framework and executive summary are available for purchase and download at `www.coso.org`.

ISO 31000

ISO 31000 ("Risk Management—Guidelines") is an international risk management framework that provides a common approach for managing any type of risk in any-sized organization and in any industry sector. The standard outlines fundamental principles and processes for contextualizing and managing risk, as well as monitoring, reviewing, and continually improving risk management procedures. The standard is available for purchase at `www.iso.org`.

Three Lines of Defense

The *three lines of defense* model is an industry practice for managing and operationalizing enterprise risk. The model defines three distinct, independent roles to ensure business operations are managed properly. The three roles are operational management, risk management and compliance, and internal audit. These roles are implemented at different levels of the organization, across business processes, as part of the wider enterprise risk management program. The personnel responsible for these three lines of defense roles are the risk owner, who is ultimately accountable and responsible for how risk is treated, and the risk practitioner, who is responsible for the day-to-day risk management functions.

The roles in the three lines of defense must be independent of one another, so that none of the roles can organizationally influence any other. Often, internal audit reports directly to the board of directors, while operations and risk are likewise separate from each other.

Operational Management

Operational management refers to the tactical management activities that take place at the business unit or business process level. At this level, risk and controls are managed on a day-to-day basis by monitoring control effectiveness, as well as any risk associated with the business process itself or its supporting systems. This line of defense can be implemented by the risk practitioner, IT, or cybersecurity teams. Operations is the layer in which controls are operated as a part of business processes and often implemented as a part of information systems.

Risk and Compliance

Risk and compliance management is a line of defense that involves proactively monitoring and managing risks associated with IT support systems, security policies that define expected behavior, security controls implemented to protect systems, and risks associated with overall business processes. This is where the organization's risk management program excels; implementing the activities associated with the risk governance framework is part of this critical line of defense. This includes overall risk assessment and analysis, risk response, and continuous risk monitoring.

Compliance management involves ensuring that the organization, its business processes, and the supporting IT systems comply with any governance requirements, whether they come from inside the organization or from external entities. Compliance directly leads to auditing processes and systems, discussed next.

While risk and compliance management are connected, they are separate and distinct entities. Risk can be variable in nature and changes over time, whereas compliance is typically a binary (yes/no, compliant/noncompliant) mentality. However, there is usually some overlap between these two areas, particularly in control assessments and risk analysis, as both risk management and compliance management often require control and risk assessments. Business process owners, senior managers, and senior risk practitioners are typically involved in this line of defense.

Internal Audit and Accountability

In the world of risk management, auditing involves continuously monitoring risk levels to ensure they remain within acceptable levels of risk appetite and tolerance, as well as monitoring the controls, management processes, and responses that maintain acceptable risk levels. This could mean auditing at the process level (user account management, for instance) or even at the overall organizational risk level, and every area of interest in between. Auditing ensures that processes and activities are being conducted in accordance with governance and serves to further ensure accountability.

While auditing and accountability are not the same thing, they are related. Accountability, in the context of risk management, ensures that risk owners are both responsible and accountable for managing risk at an acceptable level within their designated area of responsibility. In short, accountability ensures they are doing their job. Accountability is connected to auditing in the respect that only through auditing processes and activities can you determine if those processes and activities are fulfilling the governance and risk management requirements that risk owners are responsible for meeting.

Risk Profile

A *risk profile* encompasses an overall asset, organization, or business process under review. It includes detailed information on all aspects of those items and how they contribute to, mitigate, or influence risk. The risk profile for a system, for instance, would present detailed information on all the different characteristics of the system, its security controls, the risk assessment and analysis for the system, the risk responses for the system, and ongoing management for that risk. A risk profile typically consists of the following:

- Description of the asset, system, or process.
- Identified threats and vulnerabilities.
- Control inventory and effectiveness ratings.
- Likelihood and impact assessments.
- Residual risk and ownership assignment.
- Relevant regulatory or contractual obligations.

Risk profiles are dynamic: they change based on a variety of factors, including the criticality or sensitivity of the system, changes to business processes, the vulnerabilities inherent to the system, the threats that may exploit those vulnerabilities, the mitigating security controls that protect the system, and significant incidents. Regulations and societal norms also influence risk profile, particularly in areas such as data privacy. The risk profile is used to actively manage ongoing risk and ensure that it remains within acceptable levels.

Risk Appetite and Risk Tolerance

As discussed previously, the culture of the organization really defines how the organization feels about risk and how it treats risk over time. *Risk capacity* refers to the amount of loss an organization can incur without significantly affecting its ability to continue. More resilient

organizations can endure greater loss than those that are less resilient. An organization's resilience can sometimes be informed by its *risk culture*, which is senior management's attitude about risk. As part of the organization's risk culture, there are its risk appetite and risk tolerance. These are different terms you also need to know to understand risk.

Risk appetite is, in effect, how much risk an organization is willing to deal with in any given endeavor. Risk appetite is the general level of risk that an organization is willing to accept in the course of its operations. An organization's risk appetite is driven by the corporate risk culture and by the environment in which the organization exists (market, regulation, and other external factors).

Risk tolerance, on the other hand, is the acceptable level of deviation in risk for a particular endeavor or business pursuit. Risk tolerance refers to the degree of variation from the expected level of risk that an organization is willing to accept. There's a certain amount of risk in every business enterprise or pursuit; however, the organization may not be able or willing to tolerate large deviations from what it considers its acceptable level of risk on an endeavor. Note that risk appetite is often determined for the entire organization as a whole, whereas risk tolerance typically applies to a specific business venture, such as bringing a new product to market, for instance.

The board and senior management jointly define risk appetite through facilitated workshops and quantitative assessments. Stakeholder engagement ensures that appetite and tolerance reflect the organization's mission, resources, and environment. Risk appetite and tolerance must evolve in response to changes in business strategy, market conditions, or regulatory expectations. Regular reassessment ensures alignment between governance intent and operational practice.



Understand the differences between risk appetite and risk tolerance; *risk appetite* involves how much risk the organization is willing to endure, whereas *risk tolerance* is how much variation from that amount is acceptable to the business for a particular venture. Risk culture drives both of these factors.

A healthcare provider declares that it has very little appetite for risks that could expose patient health information. This strategic position reflects regulatory obligations, reputational concerns, and the high number of attacks on healthcare organizations. To ensure that operations stay within this boundary, the security team sets specific tolerance levels for system vulnerabilities.

In this organization, any system containing patient data must have all high-severity vulnerabilities patched within seven days, and no more than one medium-severity vulnerability may remain open at any time. If these limits are exceeded, the issue must be escalated immediately. The organization's low appetite establishes the strategic stance, and the tolerance thresholds guide operational decision-making to remain consistent with that stance.

Legal, Regulatory, and Contractual Requirements

Stimulus imposed by entities external to the organization typically takes the form of laws, regulations, contractual provisions, professional standards, and other requirements. Often, it is simply enough for laws and regulations to be in place, or for a company to formally accept professional standards for these requirements to be imposed on an organization. Sometimes, however, governance is also imposed as part of contractual requirements between organizations. Including governance requirements in contracts, even if an organization is not otherwise required to meet those requirements, makes them legally enforceable. In any event, organizations are required to fulfill their obligations imposed by these external governance sources.

Additionally, laws, regulations, and other governance sources often conflict with or even supplement each other. For example, one regulation may require a specific level of protection for sensitive data, and another law or regulation may require additional layers of protection based on the characteristics of the data, the organization, or the industry. For example, privacy data and healthcare data are often regulated by multiple sources of governance. In this event, the organization must implement internal policies and procedures to fulfill the requirements of all governance imposed on it. Even in cases where industry standards, such as those imposed by the Payment Card Industry Data Security Standard (PCI DSS), are in place, organizations must fulfill the requirements imposed on them if they have agreed to do so through contract or by virtue of their industry. In this example, organizations aren't legally required to comply with this standard; however, by virtue of the fact that they process credit card transactions, they are required by the credit card industry and can be censured or banned from processing transactions by the credit card companies if they do not comply with the requirements.

Compliance with legal and regulatory requirements is a critical factor in most organizations, and many of these requirements also impose risk management requirements, so risk management is often dictated, to a degree, by governance. In fact, many laws and regulations, such as the Health Insurance Portability and Accountability Act (HIPAA) and the EU AI Act, require a formalized risk management program. Risk governance frameworks, such as the NIST Risk Management Framework, also require formalized risk management. These frameworks can also be included in any other laws, regulations, contracts, or other governance requirements imposed on an organization.

Multinational organizations must be aware of data sovereignty, residency, and data transfer requirements. *Extraterritorial laws*, including the European General Data Protection Regulation (GDPR) and the EU AI Act, are laws that impose requirements on non-EU organizations that store or process data about European persons.

Professional Ethics of Risk Management

Most organizations, particularly professional ones, have requirements for a code of conduct or professional ethics, especially in the fields of cybersecurity and risk management. These codes of ethics govern the behavior of professionals, ensuring they maintain high standards of professional conduct and ethics. Most of the industry-recognized professional certifications also require adherence to a code of ethics, and ISACA is no exception.

ISACA's Code of Professional Ethics can be found at `https://www.isaca.org/code-of-professional-ethics` and applies to all organizational members and

certification holders. ISACA also implements procedures for filing complaints against professionals bound to uphold those ethical standards in the event they fail to do so. The ISACA Code of Professional Ethics (paraphrased here) includes provisions for the following:

1. Supporting and complying with standards and procedures for governance and management of information systems and technology
2. Performing duties professionally, with objectivity, due diligence, and care, as required by professional standards
3. Conducting activities in a lawful manner and maintaining the high standards of conduct and character required by the profession and ISACA
4. Ensuring privacy and confidentiality of sensitive information obtained in the course of professional duties
5. Maintaining competency in the professional field
6. Full disclosure and impartiality regarding the results of work performed to ensure that the results of that work are not distorted
7. Supporting professional education in the areas of governance and management of enterprise information systems and technology, including auditing, controls, security, and risk management

Summary

Governance guides the conduct of an organization in terms of legal and ethical behavior. It can originate from external and internal sources. External governance is imposed through laws, regulations, professional standards, risk frameworks, and contractual requirements, while internal governance arises from the organization's own policies and procedures. The organization's mission, goals, and objectives must reflect and support these governance requirements.

Senior management establishes the organization's risk appetite and tolerance. Risk capacity represents the maximum loss an organization can sustain without threatening its survival. Risk appetite refers to the amount of risk an organization is willing to accept in pursuit of its mission, while risk tolerance defines how much deviation from that appetite is acceptable for specific ventures. Mission, goals, and objectives are typically set at the strategic, long-term level.

Organizational structure influences the ability to manage risk by shaping risk culture, appetite, tolerance, and resilience. Business units and their underlying processes must address both business and IT risk at the process level. Each business process contributes to the organization's overall risk and is managed by a designated risk owner, who is accountable for the risks within their domain. An enterprise risk management (ERM) strategy must address risks at all levels—lower, middle, and higher tiers—with risk governance providing direction for the ERM program.

An organization's risk culture reflects how it perceives and responds to risk, influenced by leadership and governance. Organizational policies form part of internal governance, expressing management philosophy, culture, and leadership directives, and should align with external governance obligations.

Organizational assets encompass everything of value, including facilities, equipment, personnel, and information. Asset management includes the processes for inventorying, classifying, using, and disposing of these assets. The three lines of defense in enterprise risk management are operational management, risk and compliance management, and internal auditing.

The risk profile provides detailed information on an asset and how its characteristics contribute to, mitigate, or influence the associated risk. This profile is used to manage risk effectively. Governance may be mandated by laws, regulations, and contracts, which are legally binding, while codes of conduct and professional ethics guide the behavior of risk management professionals.

Exam Essentials

Understand how and why governance is essential. Governance represents management's control of, and visibility into, essential organizational activities. Governance consists of policies, standards, guidelines, and controls, with measurements and readouts periodically issued to management for their review.

Understand the underlying principles of the three lines model. The three lines of defense are operations, risk and compliance, and internal audit. Each must be organizationally independent of the others. Risk and compliance conducts risk assessments and establishes policies and controls that are operated by operations. Internal audit periodically examines controls and processes to ensure they are effective.

Understand the relationships between policies, standards, guidelines, processes, and procedures. Policies outline the expected behavior and characteristics of staff, as well as the processes and systems in place. Standards define how policies must be implemented and executed. Guidelines serve as non-mandatory suggestions for implementing policy. Processes are collections of one or more procedures. Procedures are step-by-step instructions for performing a task.

Recognize that asset management is a cornerstone activity essential for other processes. Asset management is essential for the integrity of several key processes, including change management, configuration management, vulnerability management, and incident response.

Know standard governance frameworks. The most well-known IT governance frameworks include ISACA Risk IT, NIST Risk Management Framework (RMF), COSO ERM, and ISO 31000.

Understand the relationships between risk appetite, risk tolerance, and risk capacity. Risk appetite is, in effect, how much risk an organization is willing to deal with in any given endeavor. Risk tolerance, on the other hand, is the acceptable level of deviation in risk for a particular endeavor or business pursuit. Risk capacity refers to the amount of loss an organization can incur without significantly impacting its ability to continue operating.

Review Questions

The following questions are designed to test your understanding of this chapter's material. For information on how to obtain additional questions, see this book's Introduction. Answers can be found at the end of this chapter.

1. An organization is subject to healthcare regulations that govern the protection requirements of individual health data. Which of the following describes this type of governance?
   - **A.** External
   - **B.** Internal
   - **C.** Regulatory
   - **D.** Professional
2. Your company handles credit card transaction processing as part of its business processes. Which of the following best describes the source and type of governance it may incur because of these business processes?
   - **A.** Internal, policies and procedures
   - **B.** External, industry standards
   - **C.** External, laws and regulations
   - **D.** Internal, industry standards
3. Your organization is structured into various departments, and each of these has its own activities that support the mission, goals, and objectives of the organization. These activities are composed at a high level and include various tasks to support the various business units. Which of the following best describes these activities?
   - **A.** Operational management
   - **B.** Organizational strategy
   - **C.** Business processes
   - **D.** Risk management functions
4. An organizational culture that is averse to risk and change is more likely to have which of the following in terms of risk appetite and tolerance?
   - **A.** High risk appetite, low risk tolerance
   - **B.** Low risk appetite, high risk tolerance
   - **C.** Low risk appetite, low risk tolerance
   - **D.** High risk appetite, high risk tolerance
5. Which of the following roles is the lowest level of responsibility in terms of risk ownership for a business process or unit?
   - **A.** Organizational CEO or president
   - **B.** VP for risk management
   - **C.** Risk practitioner
   - **D.** Business process owner

6. Which of the following is the level of variation or deviance an organization is willing to accept for a particular business venture?
   - **A.** Risk appetite
   - **B.** Risk tolerance
   - **C.** Risk acceptance
   - **D.** Risk capacity

7. A small business unit in the production department is incurring risk for one of its lower-level business processes. Although this risk is focused on the business process and not at the organizational level, it must be accounted for in the overall organizational risk assessment. At what level should this risk be considered?
   - **A.** Organizational risk
   - **B.** Operational risk
   - **C.** Strategic risk
   - **D.** Tactical risk

8. An organization needs to establish its internal risk management program, which begins with risk governance. Which of the following should the organization create first?
   - **A.** Risk strategy
   - **B.** Risk management policy
   - **C.** Risk assessment procedure
   - **D.** Risk management methodology

9. Which of the following describes why the organization exists?
   - **A.** Organizational mission statement
   - **B.** Organizational strategy
   - **C.** External governance
   - **D.** Organizational policy

10. Which of the following activities is considered foundational, without which other management activities, such as vulnerability management, incident management, and data security, could not be accomplished?
    - **A.** Configuration management
    - **B.** Risk management
    - **C.** Asset management
    - **D.** Financial management

11. A corporate policy document has a statement that reads, "For critical business processes, the organization permits no more than a low residual risk level, defined as maintaining system availability above 99.5%, preventing any unauthorized access events, and resolving identified high-severity vulnerabilities within 15 days." What type of statement is this?

A. Risk culture statement

B. Risk capacity statement

C. Risk tolerance statement

D. Risk appetite statement

12. Which of the following requires that a risk practitioner perform duties professionally, with due diligence and care, as required by professional standards?

A. Risk IT Framework

B. NIST Risk Management Framework

C. ISACA Code of Professional Ethics

D. Laws and regulations

13. A risk manager is analyzing a risk item and has determined that the organization may not survive if the risk is realized. Which of the following must the risk manager consider to make this determination?

A. Risk awareness

B. Risk capacity

C. Risk tolerance

D. Risk appetite

14. The organization's legal counsel is considering the prospect of real estate prices, including office space leasing costs, increasing significantly in the next five years. What level of risk management should be used to manage this risk?

A. Asset

B. ERM

C. Program

D. Market

15. All of the following are factors that influence an organization's culture, except:

A. Published values

B. Policies

C. Leadership behavior

D. Behavioral norms

16. Which of the following types of organizations is most likely to have a high risk appetite and a high risk tolerance?
    A. Insurance company
    B. Adult beverage producer
    C. Computer manufacturer
    D. AI agent software developer

17. How frequently should an organization review its cybersecurity policies?
    A. Annually
    B. Monthly
    C. Whenever a breach occurs
    D. Every five years

18. A new cybersecurity leader is assessing an organization's risk governance documents and records and has been unable to find any written statements on risk appetite and risk tolerance. How can historic risk appetite and risk tolerance be identified in such a case?
    A. Interview business leaders
    B. Examine incident management records
    C. Examine risk treatment records
    D. Review recent penetration test results

19. All of the following are risk management frameworks, except:
    A. NIST SP 800-37
    B. NIST SP 800-53
    C. COSO ERM
    D. ISO 31000

20. A cybersecurity consultant has observed an organization where its cybersecurity department enacts policies and controls, corporate IT implements these controls, and internal audit assesses the effectiveness of these controls. Which model has the consultant observed?
    A. Deming cycle
    B. Enterprise risk management
    C. Systems development lifecycle
    D. Three lines model

Answers to Review Questions

1. A. Laws and regulations are a form of external governance. They can also be considered as external inputs to an organization's governance program. Other examples of external inputs include market conditions, customer requirements, and supplier availability. Internal inputs include culture and policy.

2. B. This describes an external source of governance in the form of industry standards, specifically the Payment Card Industry Data Security Standard (PCI DSS).

3. C. Business processes are the activities that support the organizational mission and can be broad or broken down into detailed activities and tasks. Business processes often involve direct human activity, although IT automation enables autonomous business processes, such as the online activation of a new email account.

4. C. A risk-averse organization is likely to have both a low risk appetite and a low risk tolerance. Risk appetite is how much risk an organization is willing to deal with in any given situation, while risk tolerance is the acceptable deviation from risk appetite for a particular matter.

5. D. Although various levels of organizational management may be responsible and accountable for risk, the lowest level of risk ownership for a business process or unit is the business process owner.

6. B. Risk tolerance is the amount of variation or deviation an organization is willing to accept from its risk appetite that a particular business venture may incur. Risk appetite is the amount of risk an organization is willing to undertake. Risk capacity is the amount of loss that an organization can incur while being able to continue operations. Risk acceptance is one of four risk treatment options (the others being mitigate, transfer, and avoid).

7. D. Tactical risks are those encountered by smaller production sections—those that carry out the day-to-day work of the organization. A lower-level business process may be considered to have a tactical risk. Strategic and organizational risks are broad and may affect multiple departments or business units. Operational risks are related to ongoing, internal business operations.

8. A. The risk strategy for the organization should be created first because it provides long-term direction for the risk management program. The risk management strategy also supports external governance. Policies and procedures are subsequently created to support the strategy and implement external governance.

9. A. The organization develops an organizational mission statement to describe its overall purpose, which is its primary reason for existence.

10. C. Asset management is one of the foundational activities on which other activities depend (vulnerability management, configuration management, and so on) because assets must be inventoried and accounted for first.

11. C. A typical risk tolerance statement describes context-specific parameters that represent thresholds that would trigger escalations if exceeded.

12. C. The ISACA Code of Professional Ethics establishes requirements for ethical conduct and behavior that all certified professionals must adhere to, including the requirement to perform all duties professionally, with due diligence and care, as required by professional standards.

13. B. Risk capacity refers to the amount of loss an organization can incur without jeopardizing its ongoing operations.

14. B. A risk matter, such as macroeconomic risk, is generally managed by an organization's enterprise risk management (ERM) program.

15. B. Of the available choices, an organization's policies are the least likely to influence an organization's culture.

16. D. Of the available choices, the AI agent software developer is most likely to have a high risk appetite and high risk tolerance. The other organizations have more regulations and market reputation to consider, which reduces their risk appetite and tolerance.

17. A. The best practice for reviewing policies, including cybersecurity policies, is annually. Organizations should also review their policies when relevant regulations or technologies change, or when business conditions change.

18. C. In the absence of a formal, written risk appetite and risk tolerance, the best way to discern historic risk appetite and tolerance is to review the past few years' worth of risk treatment records. This will help the cyber leader better understand what management was willing to accept, versus what management required be mitigated, transferred, or avoided.

19. B. NIST SP 800-53 is a controls framework for U.S. government agencies and contractor organizations. The others are risk management frameworks.

20. D. The consultant observed the three lines model, also known as the three lines of defense model, in which key responsibilities are divided among three departments in a checks-and-balances model. The Deming Cycle is a lifecycle process. Enterprise risk management is those procedures and records concerning organization-wide risks of all types. The systems development lifecycle is a set of processes to implement a new system.



Chapter 2

Risk Assessment

This chapter covers CRISC Domain 2, "Risk Assessment," including the IT Risk Identification and IT Risk Analysis and Evaluation subdomains. The job practice for this domain is described by ISACA as follows:

This domain will certify your knowledge of threats and vulnerabilities to the organization's people, processes, and technology as well as the likelihood and impact of threats, vulnerabilities, and risk scenarios.

The job practice consists of these elements:

✔ A—IT RISK IDENTIFICATION

1. *Risk Events (e.g., contributing conditions, loss result)*
2. *Threat Modelling and Threat Landscape*
3. *Vulnerability and Control Deficiency Analysis (e.g., root cause analysis)*
4. *Risk Scenario Development*

✔ B—IT RISK ANALYSIS AND EVALUATION

1. *Risk Assessment Concepts, Standards, and Frameworks*
2. *Risk Register*
3. *Risk Analysis Methodologies*
4. *Business Impact Analysis*
5. *Inherent and Residual Risk*

The CRISC Task Statements relevant to this domain focus on IT risk assessment and how it applies in this context to the organization, particularly in executive management. The CRISC Task Statements covered in this chapter are:

3. *Identify threats and vulnerabilities to the organization's people, processes, and technology.*
4. *Evaluate threats, vulnerabilities, and risk to identify IT risk scenarios.*

6. *Establish and maintain the IT risk register and incorporate it into the enterprise-wide risk profile.*
9. *Conduct a risk assessment by analyzing IT risk scenarios and determining their likelihood and impact.*
10. *Identify the current state of existing controls and evaluate their effectiveness for IT risk mitigation.*

The topics in this chapter represent 22% of the CRISC examination.

The CRISC Task Statements relevant to this domain address several key activities present in a risk management program, starting with conducting IT risk assessments and the impact of IT risks on the organization's objectives or operations. These risk assessments should identify and evaluate threats, vulnerabilities, and risks, and they should be recorded in an IT risk register. In many organizations, the IT risk register is integrated into an existing enterprise risk management program. IT risk assessments should evaluate the existence and effectiveness of IT controls, identifying the need for new or modified controls based on identified risks. This risk and control information should be reported to stakeholders to facilitate risk-driven decision-making, aligning with the organization's established risk appetite and risk tolerance. IT risk personnel should periodically evaluate new and changing technologies for changes to risk. The risk management program should be aligned with industry frameworks and standards.

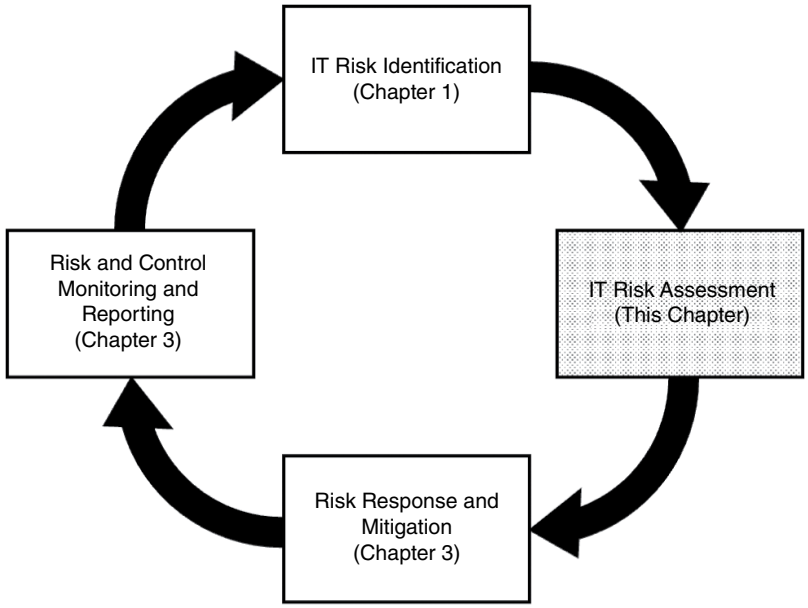
IT risk assessments are performed to identify risks associated with a specific scope of business operations or information technology in an organization. The purpose of identifying risks is for an organization to be aware of specific hazards that, if unchecked, could disrupt business operations and incur additional costs. Business executives and owners don't want bad and unexpected things to happen; rather than stick their heads in the sand, it's wiser to identify those risks and see what can be done about them. The alternative means flying blind and being unaware of risk scenarios, some of which may be highly likely to occur.

This chapter examines the comprehensive approach to IT risk assessment, which involves identifying and mitigating various risks. It also explores the detailed view of risk analysis, which involves a careful examination of individual risks, to better understand them in detail and explore various potential remedies. Conceptually, risk assessment and risk analysis are straightforward, but because there are numerous ways in which risk assessments and risk analyses can occur, along with various standards and techniques for each, all of this can become complex and cluttered.

In my opinion, ISACA's layout of the IT Risk Assessment domain isn't quite right. Both the macro (risk assessment) and micro (risk analysis) are scattered throughout the two main parts of the domain. The chapter does, however, adhere to the existing domain structure, but ensures that the context within risk management is clarified in each section.

IT risk assessments and risk analysis, covered in this chapter, are a part of the overall risk management lifecycle, depicted in Figure 2.1. Risk governance, including risk profile, risk appetite, and risk tolerance, is covered in Chapter 1, while risk treatment, emerging risks, third-party risk management, controls, and risk reporting are covered in Chapter 3.

FIGURE 2.1 The components of risk management.



IT Risk Identification

Risk identification comprises the activities performed to identify various types of risks associated with an organization's use of information technology. The types of IT risks that can be identified fall into these categories:

- **Strategic:** This category encompasses broad themes, including whether IT is aligned with the business and provides the services required by the business, as well as whether IT can support and respond to changes in business strategy.
- **Operational:** This category encompasses various issues, including IT service management, service quality, production schedules, resilience, and tactical matters such as scheduling and availability.
- **Cybersecurity:** This represents issues related to the ability of internal and external threat actors to compromise information systems or the information stored within them.
- **Privacy:** This encompasses both cybersecurity-related risks, as well as business practices related to the handling of personal information.
- **Supply chain:** This includes the availability of products such as laptops, servers, and components, as well as services like technical support, and service providers, including SaaS, PaaS, and IaaS. The supply chain itself has both operational and cybersecurity-related risks.
- **Compliance:** This represents IT-related activities that an organization may be required to perform due to a regulation, policy, or contractual terms.

Security professionals need to train themselves to recognize new risks. While at times new risks may be bold and obvious, sometimes they are subtle and easily overlooked for what they are. Often, significant risks are identified when they are least expected. In many organizations, IT risks are not apparent until a threat event occurs. The activities that may lead to the realization of new risks include the following:

- **Audits:** An audit may identify a deficiency in a business process or system that warrants more than remediation-and-forget-it treatment. Often, this requires a security analyst or risk manager to study the audit deficiency to understand its root cause. Rather than think of the deficiency as a one-time (or repeated) mistake, perhaps the design of the control is the cause, or something even broader, such as the overall design of a system or set of business processes. Where there's smoke, there's fire—sometimes.
- **Security and privacy incidents:** A security or privacy incident may reveal a latent problem that has been waiting to be discovered. The risk may not be directly related to the incident or its cause. It could be discovered by accident while an incident responder or security analyst investigates an incident more deeply.
- **Penetration tests:** While penetration tests can identify specific vulnerabilities, it is sometimes essential to read between the lines and understand the overall theme or meaning of the penetration test results.
- **News and social media articles:** Information about breaches, cybercriminal organizations, new tools and innovations, and other cybersecurity topics that can prompt a risk manager to identify risks within their organization. We have all experienced this in our careers; an example that comes to mind is the widely reported news of a critical website vulnerability at a competitor's organization. Our management contacted us to see whether we had the same or a similar exposure.
- **Security advisories:** Advisories on vulnerabilities and active attacks draw attention to specific actions and technologies. A formal cyberthreat intelligence or attack surface reduction program will subscribe to and consume all such advisories to determine whether the organization is vulnerable and then take necessary action. I consider these advisories to be more tactical than I would include in overall risk identification. However, sometimes an advisory will help us realize that something systemic isn't quite right in our organization and that we want to capture that idea while it's on our minds.
- **Networking:** Discussions about information technology, security, or privacy with professionals from other organizations can trigger the realization of a risk within one's own organization. I caution you on this point: a risk in one organization, even when similar practices or technologies are in place, does not necessarily mean that another organization faces the same risk.
- **Whistleblowers:** A disgruntled or concerned employee may choose to remain anonymous and disclose information about a condition or practice that may represent a genuine risk to the organization.

- **Passive observation:** During the course of work, a risk manager or other cybersecurity staff may identify a heretofore undiscovered risk. Sometimes a risk can hide in plain sight for an extended period, until a conversation or thought triggers a new realization.
- **Threat modeling:** This is a particular type of risk assessment, typically focused on an individual information system or business process, to identify likely threat scenarios.
- **Secure configuration assessments:** Configurations that are insecure or out of compliance can pose critical risks. This can be out of compliance with configuration baselines, such as those established by the Center for Internet Security. Others can be basic settings, such as password strength or length.
- **Governance assessments:** Governance can have a significant impact on how well the process runs. For example, if there are no status checks on systems, they may not be configured properly, thereby increasing the risk to the organization.
- **Data governance risks:** For some organizations, how data is used or shared can be very risky. For example, if a contract explicitly forbids the sharing of data but it is shared anyway, there can be legal repercussions. Reviewing the data governance process is a crucial step in protecting data.
- **Privacy assessments:** For some companies, privacy is becoming a greater concern due to enhanced regulation. Being unaware of the laws can have a significant impact on data privacy and may lead to fines and lawsuits.
- **Risk-aware culture:** The extent to which leaders, management, and staff in an organization are trained to recognize risk will result in better outcomes, including more effective security and privacy by design, and a willingness to accept that risks are always present and can be managed. Security awareness training should include risk awareness.
- **Risk assessments:** This is the wide gate through which many risks will be identified and analyzed. I place this last on the list for a few reasons: risk assessments generally have a specified scope and are aligned with a specific framework of controls, so risk analysts, by design, are looking only for particular types of risks in specific areas. These blinders can sometimes prevent someone from seeing certain types of risks, or risks in specific contexts or places.

My hope is that you will realize that the discovery of risk can occur in many ways. The best tool a risk manager can possess is an open mind, which allows them to explore the possibilities of IT risk within the organization.



The context of IT risks includes on-premises, cloud, and service-provider environments.

Brainstorming and Risk Identification

Brainstorming a concept or idea in front of a whiteboard with other colleagues is one of my favorite activities. David, a cybersecurity leader, introduces his team to a brainstorming session, where participants each have whiteboard pens and write down ideas that come to mind on a specific topic. Alternatively, one person writes down ideas spoken by others in the room.

The rule of brainstorming is this: silly ideas are not to be rejected, because that silly idea could lead to a great idea. For example, cybersecurity analysts are brainstorming ideas on zero-trust architecture. Let's join the discussion.

"Least privilege."

"Need to know."

"Multifactor authentication."

"Network access control."

"Mandatory vacations."

Wait a minute. Mandatory vacations don't seem to relate to zero-trust. But they write it down anyway. This prompts another to say, "Let's temporarily lock user accounts when they are on vacation for one week or more."

Rejecting "mandatory vacations" might not have led the team to the solution they ultimately chose, which was a bit different from what they all expected: a comprehensive zero-trust architecture.

Thoughts and discussions about risk are no different. Dare to think outside the box and consider other ideas, because this could lead you to new discoveries.

Risk Events

Risk events are the specific scenarios that, if they occurred, would inflict some impact on an organization's ongoing operations. A risk event, formally known as an *incident*, is defined as an unplanned event that has occurred or is occurring and that causes, or has the potential to cause, harm, loss, disruption, or degradation of objectives, assets, operations, or services. An incident includes a description of one or more threat scenarios, the identification of the types of threat actors, statements of vulnerabilities that could enable a threat scenario to occur, the value of relevant assets, and the probability of occurrence.

The Vocabulary of Risk

Every profession has its own lexicon—big words that represent complex topics. We aren't experts in every field, so for each of us, sometimes we overhear conversations on some professional topic that are sprinkled with words whose meaning we may not fully comprehend. Our profession, cybersecurity risk management, is no different in this regard.

Those of us in the industry occasionally hear "lay" persons misuse these words. For example, we may overhear someone in IT say, "That missing patch is a serious threat." A risk manager might hear that and think to themselves, "That person is using the term *threat* incorrectly." However, we should not fault them for trying.

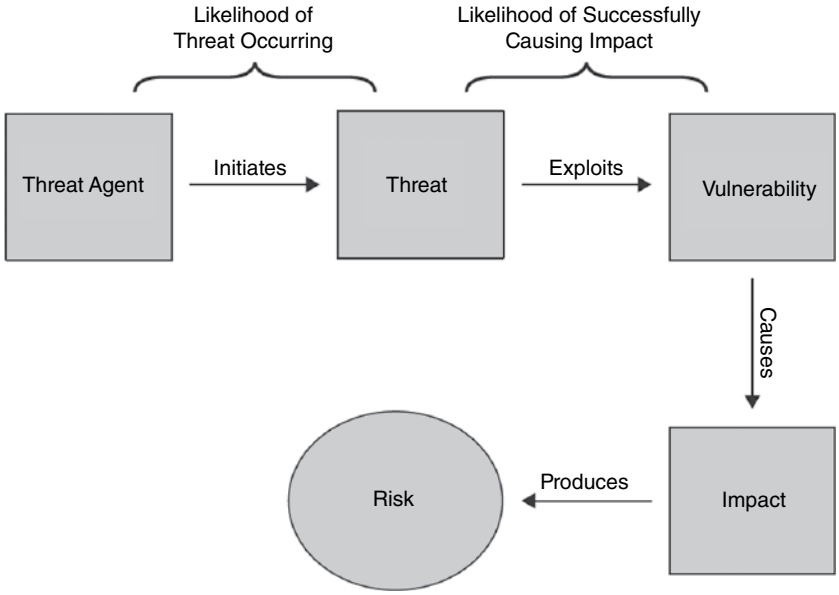
Here are the primary terms in the profession:

- ***Assessment*:** An examination of a system or process, to determine whether it is operating correctly.
- **Asset:** A tangible or intangible possession, such as a server, device, program, database, other data, intellectual property, or equipment.
- ***Vulnerability*:** A specific weakness in a system, process, or person that could fail to prevent some unwanted event. An example of a vulnerability is an easily picked lock on a door.
- ***Threat*:** An event, should it occur, that would bring harm to some asset. An example of a threat is a person who wants to break into a building and has lock-picking tools in their possession (and the skills to use them).
- ***Impact*:** The result on an organization should a threat occur.
- ***Attack*:** An offensive action that results in potential harm to an asset.
- ***Risk*:** The possibility and impact of harm that could occur to an asset, based on the asset's value, and any related vulnerabilities and threats.

Here are some additional terms that are often used:

- ***Threat actor*:** The person or entity that would carry out a threat if they chose to. Examples of threat actors include a lone hacker and a military-based cyberhacking unit.
- ***Threat agent*:** A means or mechanism through which a threat is realized. Think of a threat actor as the person, and a threat agent as the tool used by the threat actor to carry out a threat.
- ***Threat event*:** A threat that is carried out and causes some harm to an asset.
- ***Threat modeling*:** A specific type of analysis performed to better understand likely attack scenarios.
- ***False positive*:** a situation where a risk is reported that does not exist.
- ***False negative*:** a situation where a risk exists but is not reported.
- ***Exploit* (noun):** A specific tool or method that can be used to carry out a threat.
- ***Exploit* (verb):** The act of attacking an asset, made easier by the presence of a vulnerability.
- **Inherent risk:** The risk associated with a specific type of activity, before any controls are considered.
- ***Probability* (also known as *likelihood*):** The chance that a specific threat event would occur, generally within a period such as one year.
- **Residual risk:** The risk that remains after any risk reduction has been applied.
- **Risk register (also known as a risk ledger):** A listing of risks that have been identified.
- ***Risk analysis*:** The examination of one or more specific risks to determine their potential severity and impact on the organization.

FIGURE 2.2 Threats, vulnerabilities, likelihood, and impact.



- ***Risk assessment*:** An examination of a process or system to identify and describe one or more risks.
- ***Risk treatment*:** A decision made about potential action to be taken concerning a specific risk.
- ***Risk management*:** The lifecycle process that includes risk assessments, risk analysis, and risk treatment.

These are just the basic terms used in risk management, defined briefly here. These terms, and several others, are explained in detail throughout this chapter. Figure 2.2 depicts threat agent, threat, vulnerability, impact, and risk concepts.

Here are some statements you might overhear in a conversation with IT risk managers:

- "A threat actor could exploit that vulnerability on Internet-facing assets, resulting in a serious incident."
- "Threat modeling revealed additional vulnerabilities in the new system."
- "The risk assessment identified new risks that have been entered into the risk register."

Types of Risk Events

Several types of risk events can occur that are related to an organization's uses of information technology. The potential range of risk events includes scenarios such as the following:

- Breaches
- Equipment failure

- Extortion
- Natural disasters
- Fires and floods
- Sabotage
- Intrusion, data theft, and ransomware
- New regulations
- Regulatory fines
- Staff turnover
- Supplier shortages
- Strikes, riots, and demonstrations
- Terrorist attacks and acts of war

Threat Modeling and Threat Landscape

A common approach to assessing risk is to consider the universe of relevant threats that are reasonably likely to occur. Such an occurrence may compromise information or information systems, thus impacting business processes. Threat assessments and threat modeling are the primary activities used to identify relevant threats and how they may harm systems or processes.

Threat Modeling

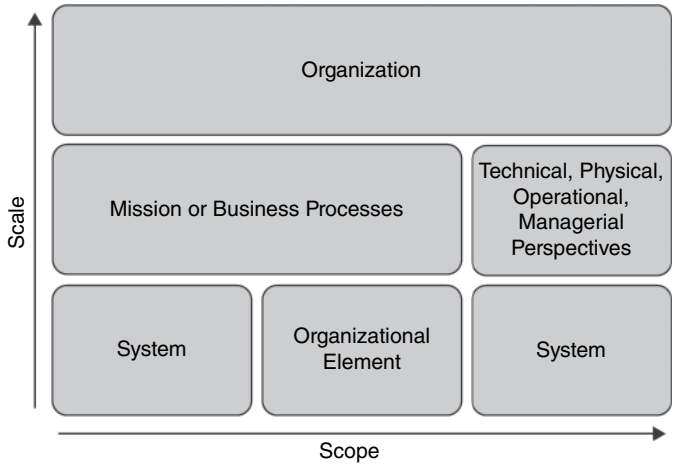
Threat modeling is a risk-identification technique that involves examining every possible threat agent, action, event, attack vector, and vulnerability for a given system, asset, or process, and then modeling or simulating how it could progress and the damage that could occur. Threat modeling has its origins in the U.S. and UK defense industries.

A *threat assessment* examines how these threats could affect the particular asset, organization, or system you are looking at, in context, and it can be done on one of several levels, either simultaneously or separately. For example, you could perform a threat assessment on a new system being developed or installed. You would examine the different threat agents and threats that could affect that particular system. You could also examine an organizational process or the entire organization itself and conduct a comprehensive threat assessment.

In addition to *scaling* a threat assessment, as described, you can examine threats from a particular perspective, performing a threat assessment specifically for technical threats, physical ones, or even external operating environment threats. This relates to the *scope* of the assessment. For example, you could look at the technical aspects of a given system or examine threats inherent to the processes and procedures associated with a particular business unit. Figure 2.3 shows this scoping and scaling process.

As Figure 2.3 illustrates, you could scale an assessment to include only specific systems within some functional regions. You could scope it by examining not only all the threats that apply to the technical aspects of systems in those particular functional areas, but also threats that affect associated business or management processes. You could have any combination of scope and scale that the organization needs to fulfill its assessment goals.

FIGURE 2.3 Scoping and scaling a threat assessment.



Threat Landscape

Threat landscape is a term that connotes a visual metaphor for considering all the threats that have been examined and identified as relevant to a given system or process. A risk manager could be overheard saying, "The threat landscape for our online commerce system is particularly hostile," meaning there are numerous significant threats that need to be considered when designing controls and defenses.

Table 2.1 illustrates how *threat agents* and threats can be categorized based on factors such as intent, relationships, skills, and others. Note that these could affect any number or combination of elements (such as processes, systems, and so on) within an organization. You could categorize threats and threat agents using these characteristics and develop others that frame and define threats and threat agents for your organization.

Note that this information is only a sampling of what you may come up with in a threat assessment. You may also break down some of these elements into far more specific ones, based on your needs and the organization and systems you are focused on. For example, a defense contractor or military organization would likely have a more specific set of threats and agents in its assessment than a nonprofit, charitable organization.



Appendices D and E in the National Institute of Standards and Technology (NIST) Special Publication 800-30, "Guide for Conducting Risk Assessments" (located on the NIST CRC website at `csrc.nist.gov`) offer a more comprehensive list of threat agents and their related threats.

Threat modeling is sometimes easier when preceded by a vulnerability assessment. You may want to conduct a threat assessment to identify the existing threats and threat agents that could exploit those vulnerabilities. These threats can be internal or external,

TABLE 2.1 Characteristics of threat agents.

Threat Category	Threat Description
Intent	**Accidental:** Untrained or careless employee and business partner **Intentional:** Hacktivist, competitor, disgruntled employee, foreign agent, unbalanced individual, organized crime, terrorist, thief, and script kiddie
Relationship to organization	**Internal:** Employees, managers, and executives **External:** Contractor, supplier, vendor, government agency, nation-state, competitor, and civil activist
Skills	None, basic, intermediate, and advanced
Objectives	**Accidental:** Workarounds or cutting corners as well as faster/less work **Intentional:** Theft, destruction, sabotage, embarrassment, blackmail, monetary gain, competitive advantage, and intelligence
Methods	**Accidental:** Inattentiveness, ignorance, arrogance, complacency, and poor planning **Intentional:** Social engineering, hacking, theft, physical damage, data alteration, and impersonation
Results	System/data theft, modification, or destruction as well as unauthorized access, compromised information, loss of business, loss of customer confidence, loss of revenue or reputation, and facility damage

human-made, or natural, and they should match a discovered vulnerability. To complete the development of a given risk scenario, you should also consider the time and location characteristics of a potential scenario. Remember that you are developing scenarios down to the most practical level of detail possible; you will have developed many different possible scenarios for an asset, each covering a particular threat–vulnerability pairing by the time you are done. You should then look at the existing controls you have in place to determine how effective they are in protecting an asset.

Organizations considering the adoption of new technology should include threat modeling in their analysis as part of the evaluation process. Threat modeling will provide the organization with insight into the scenarios that could occur if proper safeguards and controls are not implemented.



Threat and vulnerability assessments are extensively discussed in this chapter. These two assessment processes are crucial for understanding the elements that contribute to the risk associated with an asset.

Vulnerability and Control Deficiency Analysis

Vulnerability analysis and *control deficiency analysis* are important parts of risk analysis. These focus on the weaknesses in a system or process that could result in a potentially serious incident if exploited by a threat actor. The basic question that a risk analyst would ask about a given asset is, "What weakness in this system or process could be exploited by a threat actor in a successful attack?"

When stated this way, a risk analyst might think only of human-initiated incidents and overlook a broad class of vulnerabilities that are not associated with a human threat actor. Examples of non-human-initiated threats include the following:

- Lack of off-site backup media makes disaster recovery difficult in the event of a fire or flood that damages both on-site systems and backup media.
- Excessive current load on power supply circuits could result in overtemperature or power failure.

Vulnerability Analysis

Vulnerability analysis is a detailed examination of vulnerabilities that may exist or have been identified as existing. A *vulnerability assessment* can give you information about the weaknesses inherent in one or more assets. The vulnerability assessment might include examining all aspects of an asset or capability to determine any vulnerability it has, including the physical, technical, and operational vulnerabilities inherent to the asset. For example, technical vulnerabilities identified by a *vulnerability scanner* can be related to configuration issues or a lack of hardening on a network device. Physical vulnerabilities could relate to the lack of physical protection for the network device. Operational vulnerabilities might include the lack of backup processes or configuration control exercised for the device.

Control Deficiency Analysis

Control deficiency analysis is essential to understanding how controls fit into the overall risk analysis process. You perform a control analysis before determining risk, as existing controls that are already in place may be effective to varying degrees in mitigating risk. Simply examining threats, vulnerabilities, assets, impact, and likelihood without considering existing controls and protections in place doesn't do justice to the risk analysis process. Failing to account for existing controls would dramatically skew risk assessments and render an unrealistic risk determination. The goal here is to examine existing controls, determine their effectiveness, and use that determination to analyze the impact on assets if threats exploit vulnerabilities and the likelihood of such exploitation. Typically, results of your documentation, observation, interviews, and testing will likely also be part of your control analysis.

The following sections examine various aspects of control analysis. You also receive some more in-depth information on controls in Chapter 3, focusing on various aspects of how they are designed, implemented, and used to respond to risk.

Evaluating Existing Control Effectiveness

As stated earlier, risk analysis often involves evaluating the effectiveness of existing controls. This is because risk doesn't exist in a vacuum filled with only threats, vulnerabilities, and assets. Almost all assets have controls that serve to protect them from harm. These controls may or may not be sufficient to protect against particular threats. In some cases, however, a control may be implemented that significantly lowers the risk to an asset by decreasing the likelihood of a threat exploiting a vulnerability or by protecting the asset sufficiently to reduce the impact in the event a threat is realized. This is part of the risk assessment and analysis process, which involves identifying and prioritizing assets, performing threat and vulnerability assessments, and determining what controls exist and how effective they are in protecting against identified risks.

Controls fall into broad categories: administrative, technical, and physical or operational. Controls can also be categorized based on function, classifying them as preventive, detective, and deterrent. When identifying controls that apply to a given resource, process, asset, or system, you should look at all these different categories to determine what controls exist. Administrative controls will be in the form of policies or procedures and can protect an asset by detailing the requirements that users must abide by. Technical controls can protect assets by providing logical protections, including strong authentication, restrictive permissions, and other access controls. Physical and operational controls can provide asset protection through operational procedures and physical barriers, which separate unauthorized persons from the asset and prevent them from physically accessing it. Consider all these categories of controls when attempting to identify any controls that may be protecting assets.

Once you have identified all the controls protecting a given asset, you should determine what functions they are required to serve and how well they're performing those functions. Controls are designed to perform different functions (in some cases, only one function, but often many functions). Economy of use is a critical aspect of applying controls to assets; organizations that can use fewer controls to cover more security functions can reduce costs and more effectively use the resources they have.

Just as threats and vulnerabilities are often paired, controls are often identified to counter specific threats or vulnerabilities. While the value of performing threat and vulnerability assessments is discussed, seemingly as different activities, often they are conducted simultaneously, along with control analysis, to get a better idea of how these three elements balance with each other. Threats exploit vulnerabilities; however, controls can reduce the likelihood or impact of threat realization and mitigate the weaknesses that vulnerabilities provide. So, there is a relationship among these three elements, making it valuable to sometimes look at all three of them simultaneously in a concerted effort.

When assessing control effectiveness, examine how well a particular control protects against a specific threat or reduces a particular vulnerability. This can be a subjective process, but it can be more objective if performed as part of an in-depth penetration test (discussed later in this chapter in the section, "Risk Assessment Concepts, Standards, and Frameworks"). For example, a threat assessment may give you only the potential threat events that could affect a given asset, and a vulnerability assessment may tell you only what potential vulnerabilities

exist on the system. Neither tells you how effective a control may be in reducing the effects of the threat or vulnerability unless you attempt to exploit the vulnerability during penetration testing. However, by also understanding the *Common Vulnerability Scoring System (CVSS)* score associated with a vulnerability, you can better comprehend the effort required to perform an attack and the potential harm that can result from it. Only then will you truly know how effective the control is; if a given vulnerability isn't easily exploitable, then, realistically, the control may be more effective than a cursory review would indicate. Likewise, a penetration test may reveal that a control is not as effective as initially thought. Vulnerability assessments describe theoretical weaknesses, whereas penetration tests attempt to exploit vulnerabilities and give more realistic estimates of how effective controls are.

When evaluating controls as part of a risk analysis, you may categorize them as ineffective, partially effective, or effective in reducing the impact or likelihood of an event. You also might use a semi-quantitative scale that indicates control effectiveness from a range of numerical values. Each of the various control frameworks often provides guidance on assessing the effectiveness of the controls in that particular framework. For example, the catalog of controls found in NIST Special Publication 800-53, "Security and Privacy Controls for Federal Information Systems and Organizations," is supplemented by an assessment methodology found in its companion volume, Special Publication 800-53A, "Assessing Security and Privacy Controls in Federal Information Systems and Organizations," which describes how to assess and report the effectiveness of each control. Figure 2.4 shows an example of an NIST control from SP 800-53 R5, and Figure 2.5 shows the corresponding assessment guidance for the same control in SP 800-53A R5.

FIGURE 2.4 The AC-5 control from NIST SP 800-53 R5.

AC-5 SEPARATION OF DUTIES

Control:

a. Identify and document [*Assignment: organization-defined duties of individuals requiring separation*]; and

b. Define system access authorizations to support separation of duties.

Discussion: Separation of duties addresses the potential for abuse of authorized privileges and helps to reduce the risk of malevolent activity without collusion. Separation of duties includes dividing mission or business functions and support functions among different individuals or roles, conducting system support functions with different individuals, and ensuring that security personnel who administer access control functions do not also administer audit functions. Because separation of duty violations can span systems and application domains, organizations consider the entirety of systems and system components when developing policy on separation of duties. Separation of duties is enforced through the account management activities in AC-2, access control mechanisms in AC-3, and identity management activities in IA-2, IA-4, and IA-12.

Related Controls: AC-2, AC-3, AC-6, AU-9, CM-5, CM-11, CP-9, IA-2, IA-4, IA-5, IA-12, MA-3, MA-5, PS-2, SA-8, SA-17.

Control Enhancements: None.

References: None.

FIGURE 2.5 Assessing the AC-5 control, from NIST SP 800-53A R5.

AC-05	SEPARATION OF DUTIES	
	ASSESSMENT OBJECTIVE: *Determine if:*	
	AC-05_ODP	***duties of individuals requiring separation are defined;***
	AC-05a.	***<AC-05_ODP duties of individuals>*** are identified and documented;
	AC-05b.	system access authorizations to support separation of duties are defined.
	POTENTIAL ASSESSMENT METHODS AND OBJECTS:	
	AC-05-Examine	[SELECT FROM: Access control policy; procedures addressing divisions of responsibility and separation of duties; system configuration settings and associated documentation; list of divisions of responsibility and separation of duties; system access authorizations; system audit records; system security plan; other relevant documents or records].
	AC-05-Interview	[SELECT FROM: Organizational personnel with responsibilities for defining appropriate divisions of responsibility and separation of duties; organizational personnel with information security responsibilities; system/network administrators].
	AC-05-Test	[SELECT FROM: Automated mechanisms implementing separation of duties policy].



Control analysis involves identifying controls, determining their required function, assessing whether they are effective in fulfilling that function, and identifying gaps between the desired and actual states of how well controls protect assets.

Control Deficiencies and Gaps

A *control deficiency* is a situation where a *control* (including *security controls*) is not sufficient to meet its *control objective*. A *control gap* is a situation where controls are not in place. In these situations, once *control effectiveness* is determined, the organization must remediate the problem. I won't delve into remediation here; that's the subject of the Risk Response and Reporting domain, which you can read about in Chapter 3. However, during your analysis, you must identify control gaps and deficiencies and make recommendations for possible response options that would close those gaps. You may identify gaps in physical, technical, or administrative controls, even when other types of controls perform some security and protection functions.

For example, suppose an organization implements account management processes, such as creating accounts in a standardized fashion, verifying identities before user accounts are created, and so on. In that case, the organization has operational and technical controls in place that help provide secure account management. However, suppose the organization has

no written policy that requires this process and dictates its consistent performance. How can the organization ensure that these procedures are always followed uniformly and securely? If there is no policy, the organization leaves it up to individual account operators to perform this function as they want, regardless of how "most people do it." In this case, a control gap would be that, although the organization is applying some level of control over this process, there is no written policy ensuring that it is standardized.

Control Recommendations

Once you've identified existing controls, determined how effective they are, and identified gaps between the current and desired end states of controls, you should be able to make sound, risk-based recommendations to senior managers on how controls should be supplemented, added, or changed. Recommending controls is a risk-based process; remember that there is a risk you are trying to minimize, but you must balance this against the control functionality and the economy of resources. You have several considerations to keep in mind when making control recommendations:

- What more could be done to supplement or change existing controls to close the gap between current and desired end states of risk?
- Should new or different controls be applied to replace existing ones?
- What are the costs involved in applying additional controls, and do they exceed the value of the asset or the costs to replace it?
- Would additional controls require retooling or significantly changing the infrastructure, processes, or asset itself?
- Would additional controls require new personnel or training for existing personnel?
- Can additional controls be implemented that also reduce risks in other areas (economy of use)?

The answers to some of these questions will determine how you make recommendations to management about closing the "control gap." Both the cost of implementation and the *return on investment (ROI)* received from implementing new controls or strengthening existing ones will affect management's acceptance of the recommendations. The asset value, or impact if it were significantly damaged or lost, will also affect how well the recommendation is received. Implementing a costly control to protect an asset worth less in terms of impact is not usually an economically sound decision. Organizational levels of risk appetite and tolerance will also influence control recommendations, as any new or additional control must reduce risk sufficiently to fall within these levels. In general, when recommending ways to close the control gap, whether it is by introducing new controls, modifying existing ones, or modifying the infrastructure to reduce likelihood or impact, you should do the following:

- Try to leverage existing controls that can be used across the board to include additional assets.
- Look for "quick wins" first—controls that are quickly and inexpensively implemented (such as policies, training, procedure changes, and so on).
- Prioritize control recommendations based on risk; the greater the risk, the more attention to those particular control gaps.

- Be realistic in your recommendations. You can't reduce risk to absolute zero, you don't have all the resources you would like at your disposal, and you don't have to offer a 100% solution. Sometimes a 70% solution is better than a 0% solution.
- Provide alternatives to your primary recommendation for each control gap. Give management alternatives based on cost, level of risk reduction, and effectiveness.

Considering Risk and Control Ownership

Risk and *control ownership* are key considerations during a risk analysis. Depending on how the organization is structured and the risk management strategy, risk ownership may be assigned to one or several managers, spanning multiple functional areas. This is because the risk that affects one area likely affects other areas, so many different people may have responsibility for affected areas and be required to deal with and respond to risk.

Control ownership should also be examined; often, controls that protect a given asset or many systems may not fall under the operational purview of the risk owner. Some controls span the entire organization and protect multiple assets rather than only a specific system. These are referred to as *common controls*, and the entity responsible for these controls is the *common controls provider*. For example, consider physical security controls. There may be a physical security officer who oversees guards, personnel security, and access to secured areas. They may also be responsible for locks, video surveillance, and physical intrusion detection systems and alarms. These controls serve to protect information systems and data assets that might be in a secure data center.

Risks may be presented to the person who is responsible for information systems in the data center. However, the controls that serve to protect them may belong to the physical security functional area. This would require some coordination between the different functional areas if controls managed by one functional area are insufficient to protect assets managed by another area. This may also affect organizational structures, budgets, and resource allocations.



Risk, asset, and control owners are not always the same person, functional area, department, or organization. You must identify these particular owners early in the assessment process and maintain careful coordination and communication between these and other relevant stakeholders within the boundaries of your authority and assessment scope. Different types of owners can result in politically sensitive issues that revolve around resourcing, responsibility, accountability, and, sometimes, blame.

Root Cause Analysis

Root cause analysis, or RCA, is a method of problem-solving that seeks to identify the root cause of an event, situation, or problem. Often, when an event that disrupts business operations occurs, an organization wants to perform an RCA to determine what happened, why it occurred, and what changes can be made to help prevent a recurrence of the incident.

An example of RCA that often receives publicity is the analysis of an airplane crash. Here, investigators carefully attempt to reconstruct the crash to determine the ultimate cause, whether it was due to the failure of a component, pilot error, or some other factor. In cybersecurity, an organization besieged by ransomware will perform an RCA to determine how and where the attack originated.

Here are the steps in a typical root cause analysis:

Step 1—Identification: The RCA begins with a concise description of the incident or situation being analyzed. This is the starting point in an RCA.

Step 2—Chronology: Using available event data, analysis of the incident or situation involves a detailed look at how the incident or situation unfolded.

Step 3—Differentiation: Here, an investigator identifies all the details related to the event or situation, to determine which details are causal (contributed to the event or situation) and which are non-causal (did not contribute to the event or situation). Investigators recognize that many events have multiple causal factors.

Step 4—Causal analysis: Each detail or factor is analyzed to determine the degree to which it contributed to the incident or situation. This often leads to the identification of the one causal factor, the root cause, that represents the initiation of the events that led to the incident being investigated.

Step 5—Identification of corrective actions: For each causal action that contributed to the incident or situation, the RCA will identify corrective actions that reduce the probability or impact of a similar incident in the future.

A simplistic approach to RCA is known as the *five whys*. This is an iterative technique in which the analyst looks at the incident or situation (or a causal factor) and asks, "Why did that occur?" In the following example, a ransomware attack encrypted numerous files on a file server:

First track: Focus on the end user's PC.

- Why? Ransomware successfully implanted itself on an end user's PC.
- Why? The PC's antivirus was not up to date.
- Why? The PC was on a new network segment not monitored by the antivirus server.
- Why? Network engineers did not inform IT security of the new network segment (a causal factor).

Second track: Focus on the file server.

- Why? The end user had global write permissions on the file server.
- Why? It is a standard practice to grant global write access to end users, which reduces the number of access requests (a causal factor).
- Why? There are not enough service desk technicians to fulfill access requests in a more restrictive environment (a causal factor).

Third track: Focus on the end user.

- Why? The end user opened an attachment containing malware in an email message.
- Why? The end user believed that the email was sent from an executive in the company.
- Why? The end user did not recognize the indicators of a phishing message.
- Why? The end user did not attend security awareness training.
- Why? Security awareness training is not required (a causal factor).
- Why? Company executives do not believe that employees should be required to spend time away from their duties in training.

In this situation, three broad root causes were identified: first, the end user's PC was on a network not being monitored by central antivirus software because of a communication breakdown; second, the end user had excessive permissions on the file server because of a practice of loose permissions; third, end users are not required to attend security awareness training. Corrective actions will likely be recommended for all three tracks.



Why FIVE Whys and Not Four or Six?

The origin of "five whys" appears to be the Toyota Production System. While not a hard rule, some problems resolve in fewer steps, and complex systems may require more. Five is a compromise between speed and rigor. Teams experienced with root cause analysis often arrive at the same conclusion: while conducting individual RCAs, it will become self-evident when sufficient analysis has been performed.

Risk Scenario Development

A *risk scenario* is a business-impacting event that consists of a threat, a threat actor, a vulnerability, and an asset. You could say that a risk scenario is a statement about a hypothetical risk event that could reasonably occur, based on known facts about threats, threat actors, and an asset's vulnerabilities. Such a description of a risk event can help risk managers and business executives better understand what would otherwise be an abstract concept.

Risk scenario development is performed during a risk assessment, after assets, threats, and vulnerabilities have been identified and tabulated. Risk scenarios are the potential real-world events identified in risk assessment data that communicate the risks that have a tangible impact on business operations.

Risk managers can approach scenario development from two perspectives: top-down and bottom-up. Both are explained in this section.



Threat actors are not only human entities but also events such as natural disasters.

Top-down Risk Scenario Development

Top-down risk scenario development associates risks with business goals and objectives. The purpose of top-down risk scenario development is to identify scenarios that, if realized, would adversely impact those business goals and objectives. Similarly, top-down risk scenario development can associate risks with key business assets. Again, top-down risk scenarios would be related to those assets. Figure 2.6 depicts the extraction of top-down risk scenarios from risk assessment results.

Risk assessment and risk analysis are logical, analytical activities. Risk scenario development, on the other hand, particularly top-down risk scenario development, is by far an intuitive, creative activity. In pop culture, risk assessment and analysis are often associated with left-brain activity, while top-down risk scenario development is usually viewed as right-brain activity.

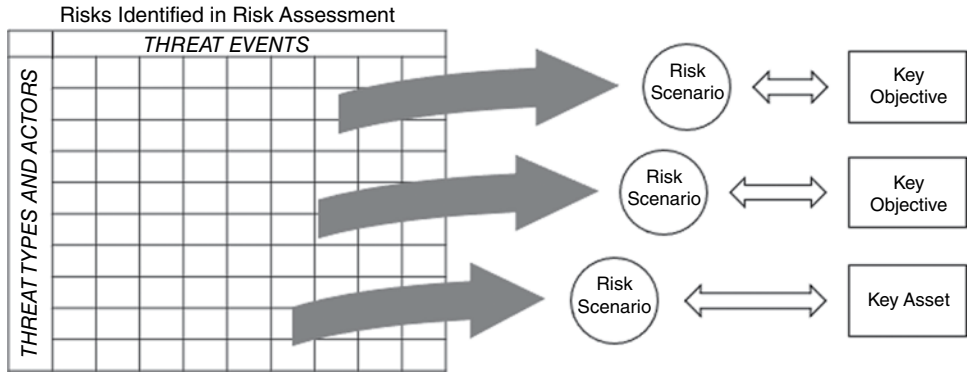
The primary benefit of using top-down risk scenarios is the ability to communicate risks more clearly to executive management. As the developers and promoters of corporate goals and objectives, senior executives should understand the various types of IT risks that could threaten to delay or increase the effort or cost required to achieve them.

Risk managers should look to outside sources to identify realistic risk scenarios. Trade journals and even mainstream media report on cybersecurity and other disruptive IT-related events. Developing a more compelling narrative than "this scenario happened to them, and it could also happen to us" can be difficult.

Bottom-up Risk Scenario Development

Bottom-up risk scenario development focuses on specific assets (or groups of assets) or threats. Bottom-up risk scenarios are more technical than top-down scenarios and less likely to be conveyed to executive management. However, this technique can help risk managers better understand the scope and nature of various types of risks. Ultimately, this may help risk managers refine their top-down risk scenario development by providing them with additional information about the universe of risks they are managing.

FIGURE 2.6 Extracting top-down risk scenarios from a risk assessment.



Risk Scenario Development Considerations

Risk scenario development is vital in risk analysis and management, where risks that may initially be abstract are refined and translated into tangible statements of risks to business operations. Absent this transformation, risk scenarios will be little more than an intellectual exercise.

To be successful, risk managers need to ensure that the risk scenarios are:

- **Relevant to the business:** Risk scenarios should indicate specific relevance and impact to the organization instead of being generic.
- **Kept current:** An organization's use of technology (at both the macro and micro levels) changes over time, meaning risks will vary in magnitude.
- **Communicated to management with a purpose:** Often, risk scenarios are the way that new risks are shared with management, allowing them to proceed with risk treatment—that is, to make decisions about how to address these new risks. This is often the only time that management is informed of risks.

IT Risk Analysis and Evaluation

After a risk has been identified (which is covered in the first half of this chapter), it can next be examined. *Risk analysis* is the detailed examination of a risk to better understand basic factors, including the likelihood and impact of risk occurrence and the development of measures that might be enacted to reduce probability or impact. Risk analysis helps determine the overall level of risk associated with a particular system, asset, business process, or even the entire organization.

Part of risk analysis involves discussing findings with risk and control owners, as well as senior managers within the organization. During these discussions, you may discover other factors that influence the likelihood and impact levels or even affect threats, vulnerabilities, and assets. You'll also review the existing controls and assess their effectiveness in mitigating risk, as well as recommend any additional controls. Finally, you will report risk analysis results to senior managers and generate a risk report after this process is complete.

These topics are covered in more detail in this section. Figure 2.7 shows a familiar variation of the 5 × 5 likelihood and impact table, demonstrating how different levels of likelihood and impact are evaluated together to produce a qualitative risk level.

FIGURE 2.7 Likelihood and impact are evaluated together to produce a risk level.

Likelihood	Impact				
	Very Low	Low	Medium	High	Very High
Very High	Very Low	Low	Medium	High	Very High
High	Very Low	Low	Medium	High	Very High
Medium	Very Low	Low	Medium	Medium	High
Low	Very Low	Low	Low	Low	Medium
Very Low	Very Low	Very Low	Very Low	Low	Low

Risk Assessment Concepts, Standards, and Frameworks

Risk assessments are performed to identify, understand, and enable organizations to make intentional decisions about individual risks. There is considerable detail within the workings of risk assessments, as discussed in this section.

Many of the terms you'll encounter on the CRISC exam are generic and may have different definitions based on the different *frameworks* or standards you are using in practice. The term *assessment* is one such term that may confuse you. For the purposes of taking the exam and in this book, an assessment refers to the overall combination of steps of identifying, evaluating, and analyzing risk. Remember that *risk identification* involves identifying the threats, vulnerabilities, assets, and existing controls. On the other hand, *evaluation* focuses on determining the likelihood and impact if the vulnerabilities are exploited by the threats, resulting in damage to the asset. Risk analysis consolidates all information, informs the organization about the risks involved, and helps determine the effectiveness of existing controls and the gaps between current and desired risk states.

Risk assessment processes sometimes include the risk identification component of the process, but are often separate from it, depending on the framework or standard being used. In general, this process involves threat source and event identification, vulnerability identification, control analysis, likelihood determination, impact analysis, risk determination, control recommendations, and results documentation. These steps follow a consistent pattern and path, as described here, for identifying risk and then evaluating it. This chapter covers some of these generic processes.



Different risk analysis methodologies and frameworks use varying terms with the same or different definitions. You should always look at a term given its context and particular framework or methodology.

Types of Assessments

Several types of assessments can be conducted to identify risks, threats, or vulnerabilities within a system or process. Some of the techniques discussed here involve different ways of thinking about risk, while others employ manual or automated tools to examine information of some type.

- **Risk assessment:** An assessment technique used to identify and classify various risks associated with systems or processes.
- ***Gap assessment*:** An assessment of processes and/or systems to determine their compliance with policies, standards, or requirements. Put another way, a gap assessment provides a view of the organization's current state, as compared to its desired future state.
- **Threat modeling:** Also known as a *threat assessment*, this is a threat-centric technique used to identify specific threat scenarios that may occur. Threat modeling is discussed in more detail earlier in this chapter.
- **Vulnerability assessment:** An assessment used to identify specific vulnerabilities that may be present in processes and/or systems.

- ***Maturity assessment***: An assessment of processes or capabilities to determine their maturity, generally according to a maturity standard such as the *Capability Maturity Model Integration (CMMI)* or the NIST Cybersecurity Framework (CSF).
- ***Penetration test***: An examination of systems, networks, or applications that identifies exploitable vulnerabilities. Tools used in a penetration test include *port scanners*, *sniffers*, *protocol analyzers*, *fuzzers*, *password crackers*, and specialized tools to identify specific vulnerabilities. A penetration test not only identifies but also validates a vulnerability by exploiting it, or by demonstrating how it can be exploited if actual exploitation would cause harm to the asset (which is particularly important when pen testing a production system).
- ***Data discovery***: A manual or automated activity in which tools or techniques are used to examine the contents of a target system to determine the presence of specific types of data. Data discovery may also involve an examination of access rights to particular types of data.
- ***Architecture review* and *design review***: A manual activity in which the architecture or design of a process or system is examined to identify potential weaknesses.
- ***Code review***: A manual activity in which software source code is examined for logic errors and security vulnerabilities.
- ***Code scan***: An automated activity in which a code-scanning program or tool examines software source code to identify vulnerabilities and other defects.
- ***Audit***: A formal inspection of a control or process to determine whether it is being followed and meets its objectives. Audits are briefly discussed in this book and thoroughly covered in the *CISA Study Guide*.

Scope of an Assessment

Before an assessment can begin, its *scope* needs to be determined. By this, I mean that the systems, processes, locations, teams, and/or time periods that are the focus of the assessment need to be formally established and communicated.

The opposing forces that influence scope are the desire to make it large enough to identify hard-to-find risks, but small enough to be efficient and complete in a reasonable period. A balance must be reached that enables the assessment to succeed.

Collecting Evidence for an Assessment

One of the challenges for a risk assessor is collecting and managing all the *evidence* for an assessment. On one end of the spectrum, the assessor may find it challenging to get the bare minimum of relevant evidence for a particular asset, including threats and threat actors, and vulnerabilities, as well as protective controls. On the other end of the same spectrum, the assessor may gather so much evidence that it is overwhelming to analyze and gauge its relevance to the assessment effort.

Another challenge is the source of the evidence. While the best evidence comes from authoritative sources with expertise in various areas, such as the asset or system in question,

current threats, and so on, the assessor must be able to separate fact from opinion and collect truly objective evidence versus subjective evidence. Objective evidence is evidence that supports an assertion about an asset or risk that can be proven or verified through independent means. Subjective evidence, on the other hand, is data whose validity is subject to opinion. For example, auditing typically produces objective evidence in the form of audit logs (assuming they are secured) that detail events occurring on a system and trace those events back to individuals, thus ensuring accountability. Events in question can be verified through the audit logs. However, system administrators who claim to regularly audit system activities but have no audit logs to support this assertion cannot have their claims verified. Additionally, if no policies or procedures require auditing or show how to perform those tasks, these events can't be independently verified either.

Evidence can come in various forms, including performance data, statistical analysis, historical data, trends, and narrative form. The assessor should establish the relevance of evidence as it is collected and determine how it relates to relevant risk questions concerning threats, vulnerabilities, likelihood, impact, assets, and controls. Overwhelming amounts of data can be reduced to relevant summaries and analyzed, sometimes in aggregate, using various automated statistical analysis and database tools.

Evidence collection typically uses a combination of four methods designed to gather various types of qualitative and quantitative data from different sources. These collection methods complement and support one another, filling in gaps from one method and verifying information collected from another. The four main evidence collection methods are reviewing documentation, conducting interviews, observing systems and assets in operation, and performing system testing on assets. They are described in more detail in the following sections.

Reviewing Documentation

Documentation is a mainstay of a risk assessment. A well-documented process or system indicates its characteristics, security requirements, and whether the people responsible for the process or system have defined measures in place to fulfill their security responsibilities. Documentation provides an assessor with a wide range of information on how the process or system is routinely used, whether a security patch program is in place, who has access to systems and data, what policies exist to define security requirements, and how the system is securely configured. Documentation can be either paper-based or electronic; it can come from spreadsheets, databases, network devices, systems, written procedures, physical access logs, and so on.

Examples of documentation include organizational policies, processes, and procedures, as well as backup and restore procedures, security officer appointment letters, and other relevant documents. While not strictly documentation (unless printed), audit logs and configuration files could also be considered in this category since they provide documentary evidence that supports an assertion of how events were recorded or how a system was configured.

Documentation also serves to verify data collected from other sources, such as interviews. Remember, if you don't document something in the security assessor world, you can't prove that you've done it.

Conducting Interviews

One of the most effective methods for collecting evidence involves interviewing key personnel, such as asset owners, system administrators, security personnel, and other people with direct knowledge of essential aspects of the asset, relevant threats and vulnerabilities, and controls. These interviews are known as *walkthroughs*. Interviews should be scheduled with key personnel and consist of relevant questions that address security posture, processes and procedures, asset characteristics, and implemented security controls. The assessor should carefully record the results of the walkthrough through concise meeting notes or recording devices for later reference. Keep in mind that while interview data should usually be considered subjective in nature, it can serve as a starting point to gather knowledge and facts about a system that can later be verified using other means, such as documentation, system observation, and testing.

The primary reason that evidence should be collected before walkthroughs is to provide the assessor with an opportunity to ask interviewees to describe processes in their own words. This practice gives the assessor an opportunity to verify whether interviewees are performing processes and procedures as described in documentation. Variations may indicate that management is not reviewing documentation sufficiently, or that process owners are not following it.

Observing Systems and Assets in Operation

Sometimes it's not enough to interview personnel and review documentation. Unfortunately, people can lie or be mistaken in interviews, and documentation can be falsified. Observing a process or system in operation can verify whether it is secure and whether its security functions and controls are working as intended to protect systems and data.

The same applies to physical security controls. An assessor can easily determine whether a control is functioning by observing a system operator perform a security function or by visually inspecting the control to ensure it is in place and properly implemented. Examples requiring the observation method include gates, guards, and locks. Other examples include watching a system through the boot process, observing a screen lock go into effect, having a system administrator demonstrate various security features, and so on.

Performing System Testing

Not all evidence can come from talking to people, reading system architecture documents, and watching the system administrator assign permissions to shared folders and files. Some evidence must come from examinations of the system's configuration, as well as relevant evidence related to patching, vulnerabilities, user accounts, configuration files, registry settings, and other technical data. This evidence is used to determine whether a system is configured properly and securely, and what potential vulnerabilities might exist on the system. This is where the technical assessment portion of evidence collection comes into play. Since technical data can be overwhelming in size and volume, automated tools are typically used to perform system security testing. These tools include vulnerability scanners, sniffers, and other hardware and software tools described later in this chapter.

During system testing, various types of host-based and network scans may be performed, along with an in-depth examination of patches and updates applied to the

system, configuration settings, user accounts, and other security settings. The system's rights, permissions, and privileges, and its resources are examined and compared against security requirements and control functions to ensure they match. Network protocols and connections are studied, and the network architecture is examined to ensure a sound network design has been implemented. Application software code reviews and testing, including fuzzing, may be performed. Authentication and encryption mechanisms are examined in depth to ensure they are robust and effective, meeting governance and best-practice requirements where applicable. Physical access controls may be tested to ensure that physical separation from systems for unauthorized users is maintained.

The results of system security assessments (usually in the form of vulnerability assessments or penetration testing, discussed later in this chapter) are consolidated with the results from walkthroughs, evidence reviews, and system observations to form a comprehensive data picture of the system or asset; these results further inform the controls analysis and impact determination aspects of risk. Table 2.2 summarizes the different data collection methods you can use during an assessment.



While you may not be directly tested on data collection methods, it is a good idea to understand these methods, as in actual practice, you will likely use them extensively.

Techniques used to examine systems or processes, evidence collection, sampling, and other considerations are covered in greater detail in ISACA's Certified Information Systems Auditor (CISA) certification and described fully in the book *CISA Study Guide* (Wiley Publishing).

TABLE 2.2 Summary of data collection methods during a risk assessment, with examples.

Method	Examples
Walkthrough	System administrators, facilities personnel, security personnel, asset managers, common control providers, risk management personnel, and control owners
Documentation review	Architectural diagrams, printed system configuration files, printed audit logs, access control logs, risk and incident reports, and security procedures
Observation	System configuration, resource access controls, physical controls, security processes and procedures as they are being performed, and system functions
Test	Systems, networks, devices, controls, procedures, ports, protocols, services, system security configuration, and vulnerability scanning

Risk Assessment Standards and Frameworks

Some organizations view risk assessment as a single, monolithic process, but many frameworks and standards break it down into distinct steps. This section examines several key frameworks, including the National Institute of Standards and Technology (NIST) SP 800-30, OCTAVE, the ISACA Risk IT Framework, ISO/IEC 27005, and others. The section goes into detail and examines them from a risk assessment perspective.

Business Alignment

One of the most important considerations for selecting standards and frameworks is business alignment. When selecting a standard for risk assessment, risk management, security controls, or integrated risk management systems, the risk management leader must choose a framework or standard that best aligns with the organization's needs. Selection considerations include the following:

- Industry sector
- Customer expectations and requirements
- Applicable regulations
- Risk appetite
- Organizational maturity
- Current and future technology in use

Security and risk professionals who strive for business alignment will achieve better outcomes by building and selecting processes, frameworks, and standards that are more likely to succeed.

NIST Special Publication 800-30

NIST Special Publication 800-30, "Guide for Conducting Risk Assessments," is a step-by-step guide to performing qualitative and quantitative risk assessments. The standard also briefly describes the overall risk management lifecycle process to provide context for individual risk assessments. The steps of a risk assessment are as follows:

1. ***Prepare for assessment.*** The scope, techniques, records, personnel, schedule, and other resources are identified. You should also consider the threat landscape, although this is not explicitly part of NIST SP 800-30.
2. ***Conduct assessment.*** The following are the actual risk assessment steps performed on each target asset defined in the assessment scope:
   - ***Identify threat sources and events*:** All reasonable threat sources, events, and actors are identified.

- ***Identify vulnerabilities and predisposing conditions*:** Vulnerabilities in the in-scope processes or systems are identified. Often, vulnerabilities and threats are paired and listed. For instance, a threat actor with a lock-pick set is paired with an exterior door with a substandard key core.
- ***Determine the likelihood of occurrence*:** The probability that each specific threat will occur is determined. Often, this is an estimate, as there is generally insufficient data to reliably determine an accurate probability.
- ***Determine the magnitude of impact*:** This is typically expressed in qualitative terms, in the form of descriptions of the impact on business operations. Also, this can be expressed in quantitative terms that might include response and recovery costs.
- ***Determine risk*:** The risk associated with each threat scenario is calculated, generally as Risk = Impact × Probability × Asset Value.



NIST 800-30 states risk as Risk = Impact × Probability, but other methodologies include asset value and process criticality, which help to further distinguish higher risks from lower risks.

3. *Communicate results.* The results of the risk assessment are tabulated as determined at the start of the assessment. The results are communicated to various stakeholders in the organization. Often, results are developed for specific audiences; for instance, an executive summary may be produced for executives, while detailed charts will be produced for lower management and those who will participate in remediation activities.
4. *Maintain assessment.* This generally refers to the activity of risk monitoring, in which a risk manager periodically re-examines the assets to determine whether any changes have occurred in threats, threat actors, vulnerabilities, impact, or probability of occurrence.

NIST SP 800-30 contains additional resources, including the following:

- **Appendix D: Threat Sources:** This includes taxonomies of adversarial and non-adversarial threat sources, capabilities, intent, and scale.
- **Appendix E: Threat Events:** This includes numerous adversarial and non-adversarial threat activities that represent reconnaissance or attack.
- **Appendix F: Vulnerabilities and Predisposing Conditions:** This is a taxonomy of vulnerabilities and their associated characteristics.
- **Appendix G: Likelihood of Occurrence:** This includes taxonomies related to the probability of threats occurring.
- **Appendix H: Impact:** This includes various types of impact on organizations when threats occur.
- **Appendix I: Risk Determination:** This includes taxonomies used to determine the magnitude of individual risks.

OCTAVE Methodology

As part of a U.S. government contract with Carnegie Mellon University, the Operationally Critical Threat, Asset, and Vulnerability Evaluation (*OCTAVE*) methodology was developed to assist organizations in identifying and assessing information security risk. The methodology was initially developed by the Software Engineering Institute (SEI) at Carnegie Mellon University in 1999 and has been updated over the years to its current version, OCTAVE Allegro, in 2007.

OCTAVE uses a concept called *workshops*, made up of members of an organization, sometimes including outside facilitators with risk management and assessment expertise. The methodology has prescribed procedures, including worksheets and information catalogs to assist organizational members, drawn from all functional areas of an organization, to frame and assess risk based on internal organizational context (infrastructure, governance, business environment, and so on). OCTAVE has three iterations: the original OCTAVE, OCTAVE-S, and OCTAVE Allegro.

OCTAVE

The original OCTAVE methodology was released in 1999 and was designed for organizations with at least 300 employees. The assumptions for using OCTAVE included that the organization operates and maintains a multitiered information infrastructure and performs vulnerability assessments. While OCTAVE is flexible and allows an organization to tailor its use based on its own assessment needs, the methodology is performed as three sequential phases. The first phase covers identifying assets, threats, protection practices, vulnerabilities, and security requirements. You can see how this aligns with the practices described throughout this chapter, including determining threats, identifying assets and their vulnerabilities, and evaluating existing controls. In the second phase, the members of the workshop perform assessment activities and evaluate the infrastructure. In the third phase, the workshop members develop response strategies for the identified risks.

OCTAVE-S

OCTAVE-S is fundamentally the same as the original OCTAVE, with a few minor differences. First, OCTAVE-S was designed to assess smaller organizations, usually with fewer than 100 people. Second, this iteration relies more on internal personnel with extensive knowledge of the organization and its infrastructure and less on outside risk experts or facilitators. The number of workshop members will also be considerably less; the team may rely on fewer than ten key experts in the organization. OCTAVE-S processes and procedure documents were also written to include more detailed security-related information. This approach does not rely on outside security experts or risk facilitators to assist the team.

OCTAVE Allegro

The OCTAVE Allegro methodology, introduced in 2007, expands the previous two iterations but does not necessarily replace them. It includes a more business-centered and operational risk approach to the assessment. It is also asset-centered in its approach;

it focuses on assets and their use or misuse, the environment they are used in, and how threats and vulnerabilities specifically target and affect these assets. Allegro can be used on a scaled basis; either individuals or workshop teams can use this method without requiring a high degree of risk management or assessment knowledge and experience. OCTAVE Allegro is divided into four phases, which include eight steps. These phases and steps are described next.

Phase 1: Establish Drivers

Step 1: Establish risk measurement criteria. This step serves to develop and establish the organization's methodology and measurement criteria used to assess risk. OCTAVE employs a qualitative assessment methodology and measurements, but can use quantitative methods for certain aspects of the overall process, such as determining likelihood and impact.

Phase 2: Profile Assets

Step 2: Develop an information asset profile. In this step, the organization develops an asset profile, which is a collection of information that describes the asset, including its characteristics, priority, impact on the organization, and value. This profile also contains security requirements the asset may have.

Step 3: Identify information asset containers. The information asset container describes how data is stored, where it is processed, and how it is transported. Containers are usually networks and systems, including those that the organization directly operates and those that it outsources.

Phase 3: Identify Threats

Step 4: Identify areas of concern. In this step, potential risk factors are identified and used to develop threat scenarios.

Step 5: Identify threat scenarios. In the context of OCTAVE, threat scenarios describe categories of actors and the relevant threats in each category. The scenarios are typically identified by using a threat tree, which simply maps actors and scenarios.

Phase 4: Identify and Mitigate Risks

Step 6: Identify risks. In this step, consequences such as impact and likelihood are identified and measured, which inform risk.

Step 7: Analyze risks. Once impact and likelihood have been identified and measured, risks are analyzed. During this step, risk scores are developed. These come from developing impact and likelihood values, emphasizing impact in particular, since this is an asset-driven assessment.

Step 8: Select mitigation approach. In this step, risk response strategies are developed, analyzed, and recommended.



OCTAVE is designed for larger organizations, while OCTAVE-S is intended for use by smaller organizations. OCTAVE Allegro is scalable for use by individuals or larger teams working in organizations of various sizes.

ISO/IEC Standards

ISO stands for the International Organization for Standardization, and *IEC* stands for the International Electrotechnical Commission. Together, these two organizations are responsible for a wide range of information technology and manufacturing standards used worldwide, including several that apply to information security and risk management. These standards are available from `www.iso.org`.

ISO/IEC 27005:2022

Standards in the *ISO/IEC 27000 series* focus on information security, and the 27005:2022 standard ("Information technology—Security techniques—Information security risk management") supports the common ideas promulgated in both ISO/IEC 27001 and 27002. This standard defines the entire risk management lifecycle, including detailed risk assessment principles. The standard doesn't offer specific risk management or assessment methodologies. However, it does serve to describe a formalized, structured risk management process, such as developing context, scope, methods, and so on. It also describes the two primary types of assessment methodologies: qualitative and quantitative. Additionally, it describes processes used to develop risk response strategies, communication with stakeholders, and continuous risk monitoring.

IEC 31010:2019

IEC 31010:2019[1] is where the meat of information regarding risk assessments is located for the ISO/IEC standards. This document, "Risk Management—Risk Assessment Techniques," articulates and extends the risk management principles of ISO 31000:2018 and provides a concrete overview of the risk assessment process. This standard describes risk assessment as the combination of risk identification, analysis, and evaluation, preceded by a step of establishing the risk context within the organization. The standard also addresses communication and consultation with various stakeholders and management, risk treatment (response), and monitoring risk.



Of the ISO/IEC standards, IEC 31010:2019 is relevant for the exam because it covers risk assessments in detail.

ISACA's Risk IT Framework

ISACA developed the Risk IT Framework to align with its *COBIT* framework. While COBIT is used to manage other aspects of business infrastructure and risk, the Risk IT Framework is explicitly used to allow organizations to identify and manage IT risk. The framework is broken into three major process areas: Risk Governance (RG), Risk

[1] This standard is reviewed and revised every five years. A newer version may be available by the time this book is published.

Evaluation (RE), and Risk Response (RR). This framework was briefly covered in Chapter 1, but the focus of this chapter is on the assessment process, so it covers the Risk Evaluation portion in more depth.

The Risk Evaluation processes in the Risk IT Framework don't cover only assessment; they overlap with the risk identification areas. In RE processes, business impact is framed and described, and risk scenarios are also developed. Risks are assessed, analyzed, and presented to the organization's management. The RE processes are further broken down into three areas (numbered RE1–RE3), which the following sections describe in more detail.

Collect Data (RE1)

The data collection process aligns with some of the risk identification processes previously described. During this process, risk management personnel develop a model for data collection, which provides for standardized data formats, measurements, and common data definitions. Data is collected on the various aspects of the organization and risk scenarios, including the business's operating environment, risk factors, threat sources and events, vulnerability data, and asset data. Data is also collected on the effectiveness of existing controls in the organization.

Analyze Risk (RE2)

In the Analyze Risk part of the Risk IT Framework, the organization begins to assess and analyze risk. First, it defines the risk analysis scope. The scope determines how broad and deep the risk analysis efforts will be, what areas the risk analysis will cover, and which assets will be examined for risk. To complete this part of the process, you'll need to consider all the documentation assembled up to this point, including risk scenarios, asset inventories, breakdowns of business processes, and prioritization of assets, among others. This will help frame the risk analysis as well as determine scope.

During this step, IT risk is also estimated. This involves determining likelihood and impact values associated with the developed risk scenarios. You will also need to consider existing controls and assess their effectiveness as currently installed and functioning. Likelihood and impact values are considered after existing controls are considered.

Although Risk Response is the subject of Chapter 3, you also identify and consider possible Risk Response options during this step of the process. The person assessing the risk may recommend several different response options, based on the identified risk. What risk options an organization will use is determined later in the process and typically with the input and approval of senior management. Risk Response options should be recommended that reduce risk to an acceptable level, directly based on the risk appetite and tolerance of the organization. Finally, knowledgeable peers should review risk analyses to verify that the analysis process is sound and to validate the results.

Maintain Risk Profile (RE3)

A *risk profile* is a collection of detailed data on identified IT risks. Risk profiles can cover a single system or asset, but are also often seen as describing risks on an organization-wide basis. During this Risk Evaluation step, a comprehensive document of identified risks

and their characteristics, such as details regarding impact, likelihood, and contributing factors, is developed and maintained throughout the asset or system's lifecycle. Remember that system or asset risk profiles are usually rolled up into a more comprehensive risk document that covers systems, assets, and business processes across the entire organization.

IT assets are also mapped to the business and organizational processes they support during this step; this helps translate IT risk to corresponding business risk and allows management to see how the organization would be affected if IT risk were realized on specific assets. This also enables the organization to develop criticality estimates of IT resources from a business perspective. It's worth noting that this process is similar to the one followed during a business impact assessment (BIA).

The key to understanding how IT risk affects business operations lies in understanding how IT capabilities are provided to the business processes in question and how critical IT is to a particular business process or operation. Maintaining a risk profile also means monitoring and updating risk scenarios and analysis as conditions and risk factors change. This involves keeping the risk register up to date as well. Finally, the organization should develop key risk indicators (KRIs) during this step, specifically focusing on IT resources. While key risk indicators are discussed in Chapter 3, for now, you should know that they allow the organization to monitor changes in risk for given scenarios, enabling the organization to modify its risk profiles as these indicators change. Table 2.3 summarizes some of the different risk assessment methodologies available to you.

TABLE 2.3 Summary of available risk assessment methods and frameworks.

Methodology	Sponsor	Notes
NIST SP 800-30, "Guide for Conducting Risk Assessments"	National Institute of Standards and Technology	Part of the NIST Risk Management Framework (NIST RMF)
OCTAVE (OCTAVE, OCTAVE-S, and OCTAVE Allegro)	Software Engineering Institute, Carnegie Mellon University	Uses workgroups, is suited for any organization, and does not require extensive risk management experience or knowledge
ISO/IEC	ISO/IEC	Prescribes methods and processes for developing assessment programs as well as describes more than 30 different means to assess risk factors and elements
The ISACA Risk IT Framework	ISACA	Risk assessments covered in the Risk Evaluation (RE) phase



Because this is an ISACA-sponsored exam, it is a good idea to understand risk assessment terms and processes from the perspective of the ISACA framework more so than any other. The ISACA Risk IT Framework also aligns with ISO/IEC standards.

Risk Ranking

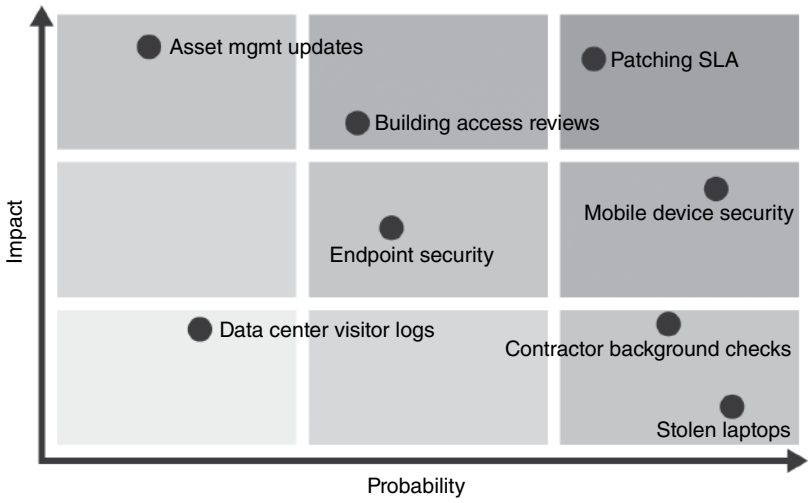
When risks have been identified in a risk assessment, it generally makes sense to display these risks in an ordered list, resulting in *risk ranking*. Most often, the highest risks will appear at the top of a list, indicating the nature of the sequencing. For instance, risks with the most significant impact may appear first, or perhaps a composite risk score will be calculated for each risk (based on probability, impact, and asset value) and risk items sorted by this overall risk calculation. Risk ranking, then, helps tell the story of the most critical risks that have been identified.

Risk ranking can be portrayed visually, which helps the reader better understand the universe of risks that have been identified. Figure 2.8 shows a risk map that visually indicates the nature and severity of risks.

Risk Ownership

To effectively manage and act on unacceptable risks, each risk is assigned a *risk owner*. Generally, the risk owner is the department head or business unit leader associated with the business process that is the focus of a risk. It is usually not appropriate to assign

FIGURE 2.8 A risk map visually depicts the severity of risks.



risk ownership to IT, as IT serves as the steward of information and information systems, while also providing services directed by department heads and business unit leaders.

Similarly, the organization's cybersecurity leader (often a CISO) is not the assignee of risks. Instead, the cybersecurity leader facilitates risk conversations and decisions with department heads and business unit leaders, particularly those whose business processes are associated with individual risks.

In risk management, there are usually several owners associated with each risk, including the following:

- **Risk owner:** This is usually the department head or business unit leader associated with the business process where the risk resides. Even if a risk is related to an IT system, IT is generally not the risk owner. Rarely, if ever, is the cybersecurity leader assigned as a risk owner.
- **Risk treatment decision-maker:** This person makes a business decision regarding the disposition (treatment) of an individual risk. Risk treatment is discussed fully in Chapter 3.
- **Risk remediation owner:** This is the person responsible for any remediation that has been selected as a part of risk treatment.

An organization's risk management program charter (or other formal document) should define the roles and responsibilities of each type of risk owner.

These and other facets of risk ownership are generally indicated in the risk register, which is discussed next.

Risk Register

A *risk register*, also known as a risk ledger, is a business record containing information about business risks and their origin, potential impact, affected assets, probability of occurrence, and treatment. A risk register is *the* central business record in an organization's risk management program (the set of activities used to identify and treat risks).

Together with other records, the risk register serves as the focal point of information that an organization is attempting to manage risk. Other records include information about risk treatment decisions and approvals, tracking projects linked to risks, risk assessments, and other activities that contribute to the risk register.

When new risks are identified, they are added to the risk register to be later analyzed, discussed, and decisions made about their disposition.

A risk register can be stored in a spreadsheet, database, or within an *integrated risk management (IRM) system*—formerly known as a governance, risk, and compliance (GRC) tool—used to manage risk and other activities in the security program. Organizations with an enterprise risk management (ERM) system will have a risk register there. Table 2.4 shows a typical risk register entry.

TABLE 2.4 Sample risk register data structure.

Item	Description
Entry number	A unique numeric value identifying the entry. This can be in the form of a date, such as 20260127a.
Status	Current status of the entry: • Open • Assigned • Closed
Date entered	The date the risk register entry was created.
Entered by	The person who created the risk register entry.
Source	The activity or event that compelled someone to create this entry. Sources include the following: • Risk assessment • Vulnerability assessment • Security incident • Threat intelligence • External party
Incident number	Reference to an incident record, if applicable.
Title	Short title describing the risk entry.
Description	Description of the risk.
Threat description	Description of the potential threat activity.
Threat actor	Description of the type of threat actor: • Worker • Former worker • Supplier, vendor, or partner • Cybercriminal • Nation-state
Vulnerability description	Description of one or more vulnerabilities that increases the probability or impact of threat realization.
Third-party organization	Name of the third-party organization where the risk is present, if applicable.
Third-party classification	Classification level of the third-party organization, if applicable.

Item	Description
Business impact	Business language description of the impact of threat realization.
Technical impact	Technical language description of the impact of threat realization, if applicable.
Asset	The specific asset, asset group, or asset class affected by the risk.
Asset owner	The owner of the affected asset.
Risk owner	The owner of the risk.
Control group	A reference to the affected control group, if applicable.
Control	A reference to the affected control, if applicable.
Process	A reference to the affected process, if applicable.
Untreated probability of occurrence	An estimate of the probability of occurrence of the threat event associated with the risk. Usually expressed as high, medium, or low or as a number on a scale such as 1 to 5.
Untreated impact of occurrence	An estimate of the impact of occurrence of the threat event associated with the risk. Usually expressed as high, medium, or low or as a number on a scale such as 1 to 5.
Untreated risk score	An overall risk score that is generally a product of probability, impact, and asset value.
Treated probability of occurrence	An estimate of the probability of occurrence of the threat event associated with the risk after risk treatment. Usually expressed as high, medium, or low or as a number on a scale such as 1 to 5.
Treated impact of occurrence	An estimate of the impact of occurrence of the threat event associated with the risk after risk treatment. Usually expressed as high, medium, or low or as a number on a scale such as 1 to 5.
Treated risk score	An overall risk score that is generally a product of probability, impact, and asset value, after risk treatment.
Estimated cost of risk treatment	An estimated cost of risk treatment. This is expressed in dollars or the local currency.

(*Continued*)

TABLE 2.4 (Continued)

Item	Description
Estimated level of effort of risk treatment	An estimated level of effort of risk treatment. This can be expressed as high, medium, or low, as a number on a scale such as 1 to 5, or as an estimate of range of man-hours, as follows: • Less than 1 hour • Less than 10 hours • Less than 100 hours • Less than 1,000 hours • Less than 10,000 hours
Risk treatment	The chosen method of risk treatment: • Accept • Mitigate • Transfer • Avoid
Risk treatment approver	The person or body that approved the risk treatment method.
Risk treatment approval date	The date that the risk treatment method was approved.
Risk treatment owner	The person responsible for carrying out risk treatment.
Risk treatment description	A description of the risk treatment.
Risk treatment planned completion	The date when risk treatment is expected to be completed.
Actual cost of risk treatment	The actual cost of risk treatment, which would be known when risk treatment has been completed. This is expressed in dollars or the local currency.
Actual level of effort of risk treatment	The actual level of effort of risk treatment, which would be known when risk treatment has been completed. This is expressed in work hours.
Risk treatment closure date	Date when risk treatment is completed.



A higher purpose of a risk register is triage. Qualitative and semi-quantitative risk ratings help an organization distinguish more significant risks from those that are less significant.

Sources of Information for the Risk Register

Awareness of risks can come from many sources and through various events. The information in Table 2.4 provides some hints about the potential sources of information that would lead to the creation of a risk register entry. These sources include the following:

- **Internal and external risk assessment:** A prime source for risk register entries, a risk assessment identifies risks within the organization's staff (for example, excessive workload, competency, and training issues), business processes, and technology.
- **Vulnerability assessment:** The high-level results of a vulnerability assessment (or penetration test, code review, social engineering assessment, etc.) may indicate overarching problems in staff, business processes, or technology at a strategic level.
- **Internal audit:** Internal audit and other internal control self-assessments can identify problems in staff, business processes, or technology.
- **External audits:** Audits performed by external parties can identify issues in business processes and information technology.
- **Security incident:** The occurrence of a security incident may reveal the presence of one or more risks that require attention. Note that a security incident in another organization may highlight risks in one's own organization.
- **Threat intelligence:** Formal and informal subscriptions or data feeds on threat intelligence may reveal risks that warrant attention.
- **Industry development:** Changes in the organization's industry sector, such as new business activities and techniques, may reveal or magnify risks that require attention.
- **New laws and regulations:** The passage of new laws, regulations, applicable standards, and private legal obligations may reveal the presence of risks that require attention. Also, note that compliance risk (i.e., the possibility that regulators or others may impose fines or sanctions on the organization) may well be included in one or more risk register entries if the organization has identified such risks.
- **Consultants:** A visit by or conversation with an expert security consultant may reveal previously unknown risks. The consultant, who may be an auditor or assessor or may be working in the organization on a project, may or may not expect to find risks that the organization's security manager would want to be aware of.

Updating the Risk Register

Because the risk register is a living document, it should periodically be updated as risk conditions change. Risk assessments will provide updated information that should be included in the risk register and information about ongoing risk response actions applied to the risk. Whether the risk register is in the form of an electronic database, spreadsheet, or another type of document, it should be easy to read and readily accessible to risk managers and risk owners, allowing them to stay updated on the organization's current risk posture. As mentioned previously, it should also be considered a sensitive, access-controlled document. Table 2.5 shows a risk register and some sample entries.

TABLE 2.5 Sample entries in a risk register.

Risk Description	Likelihood	Impact	Risk Owner	Resources Required	Estimated Completion Date	Actual Completion Date	Current Status	Comments
Data theft from a malicious user (via risk assessment in April 2025)	Medium	High	Engineering department	$10,000 for new DLP system	January 2027		Assessing different DLP vendors	Purchase and implementation by December 2026
Loss of primary accounting server data	High	High	Accounting department	$15,000 for backup system	October 2025	September 2025	Vendor selected	Final purchase and installation in August 2026

Risk Analysis Methodologies

As stated earlier in this chapter, *risk analysis* involves a detailed examination of a risk identified during a risk assessment. In risk analysis, I examine a risk and apply techniques to better understand it from several perspectives:

- **Asset identification:** The asset(s) involved in a particular risk are identified. This could be an IT or IoT device, a business application, an entire IT environment (or specific parts therein), a database, a building, a procedure, or a business process.
- **Vulnerability:** Various weaknesses related to the asset are identified that could be exploited by a threat event.
- **Threat:** Various threat scenarios are considered, particularly those that are reasonably applicable to identified vulnerabilities.
- **Asset value:** The value of the asset is identified and considered. Various iterations of risk analysis may consider different types of *asset valuation*, including replacement cost, repair cost, lost revenue, and more. For various risk scenarios, it may be necessary to use and identify different types of asset value. For instance, one risk scenario may use replacement value while another might use repair cost.
- **Business context:** Because risk analysis is often an abstract mental exercise, it is necessary to define the business context for the risk. For example, risk analysis is performed on a database server; the business context is the organization's financial accounting system that supports all of the organization's accounting department's business processes. Without identifying the business context, a risk remains abstract, and it is more difficult to apply risk treatment later on.
- **Business impact:** To be relevant to the business, risk analysis must identify the impact of various threat scenarios on business processes and operations. For example, risk analysis focused on access controls on a database management system containing customer information would identify business interruption in scenarios involving data destruction and privacy breaches, with reputation damage and fines in scenarios involving data theft.
- **Risk appetite and risk tolerance:** The risk analyst should be familiar with the organization's stated risk appetite and risk tolerance to ensure that identified risks and subsequent risk analysis align with these parameters. Risk appetite and risk tolerance are discussed in more detail in Chapter 1.
- **Mitigation analysis:** Once a specific risk is examined, various mitigation scenarios can be developed. For each, the risk manager would identify the time required to perform the mitigation, the associated cost, the changes in the probability of various threat scenarios, the changes in impact for each scenario, and the resulting change in risk.

Asset Identification, Classification, and Valuation

Identifying the different elements of risk is a necessary precursor to assessing risk. One of the goals of the identification process is to develop risk scenarios. Risk scenarios require that you identify assets, vulnerabilities, and potential threat events.

Asset identification not only requires you to inventory all of your assets; you should also determine how critical they are to your business, decide how sensitive the data they process or produce is, and characterize or classify the system or asset in terms of confidentiality, integrity, and availability. Determining criticality is an essential aspect of the overall risk management process and is typically conducted before any assessment is undertaken. The NIST RMF process requires that you perform a system categorization on assets in the first step of that framework. This system categorization later helps you to assign controls to the asset based on its criticality factors. The set of controls you assign to each asset, based on its criticality, is assessed for effectiveness in mitigating risk and ensuring compliance.

Risk managers may also want to know the value of an asset. Quantifiable risk analysis requires that some value of an asset be identified, whether based on replacement cost, salvage cost, depreciated value, or another method of calculation that represents the expenditure of resources to return an asset to its pre-threat event state.

Likelihood Determination and Impact Analysis

Evaluation is the next step of the process. You must determine the likelihood of a given threat agent initiating a threat or of an identified threat exploiting a given vulnerability, as well as the level of impact or harm to the asset if a threat has successfully exploited a vulnerability. While threats and vulnerabilities in assets are necessary ingredients to determine risk, they are static; in other words, threats, vulnerabilities, and assets don't change much. On the other hand, likelihood and impact are variable and are the primary factors in measuring risk. People who don't understand how risk is framed often mistakenly attribute risk as a result of only one individual element. You could say that the risk to an asset is high, for example, simply because the likelihood of the adverse event occurring is high. You could also say that risk is high because the level of impact if an adverse event were to occur would be high. However, neither of these statements alone really captures the sense of risk. This is because they do not convey the other dimension of risk individually, since risk is a product of *both* likelihood and impact. Additionally, you can't make these determinations without the background knowledge of the threat, the vulnerability, and the asset.



Make sure you consider controls and their effectiveness when analyzing risk, since controls can help decrease likelihood or reduce impact, thereby reducing the assessed risk for an asset and the organization.

Evaluation involves measuring two primary factors: likelihood and impact. This chapter has already discussed how there can be different ways to measure these factors based on subjectivity or concrete values. This is where the discussion on qualitative and quantitative techniques comes into play. The next section delves into those assessment techniques in more depth.



Remember that risk assessments consist of three key steps: risk identification, evaluation, and analysis. *Identification* produces threat–vulnerability pairs, asset inventories, and risk scenarios. *Evaluation* covers likelihood and impact calculations, resulting in risk values. *Analysis* looks at the overall application of these factors, as well as control analysis, to produce a comprehensive picture of risk.

Quantitative and Qualitative Analysis Techniques

If you asked two different security professionals, both trained in the latest standards and practices associated with risk management, to assess the same system, chances are that they would produce similar reports. However, they would not be the same. Why is that? It's because there are so many different methods of conducting risk assessment and analysis, depending on what governance the organization falls under, how the organization has framed its risk, and so on, that the risk would never be assessed the same; there are too many variables involved. Add to that complexity the natural subjectivity of human beings, and you may have very different risk reports. However, most risk assessment methodologies are still fundamentally based on the two primary ways to assess risk: qualitative and quantitative. Both methods have value, depending on the context of the situation, and often, both methods can be combined. The upcoming sections discuss these techniques at length and explain how they can be used together.

Quantitative Risk Analysis

Quantitative risk analysis uses concrete (nonsubjective) numerical values and statistical analysis to calculate the likelihood and impact of risk. Quantitative methods use numerical values such as dollar figures, statistical numbers, and other hardcore numerical assignments. *Quantitative assessments* follow the same general pattern described previously. You must first list your assets (along with their replacement costs and exposure factors), calculate the impact of a full or partial loss of the asset, determine the likelihood of the adverse event, and assess the overall risk to the organization. This might be calculated in terms of dollars, for example.

Let's look at a scenario to help you better understand quantitative risk analysis, including some of the factors that can be quantified and how.

Let's say you are a new employee in a small family-owned business that has no other security professionals on staff. You've been tasked with developing a risk management strategy and program for the business. Suppose the business has a small server farm located in an older building on your small campus, with wooden walls, older wiring, no backup electrical power, and reasonably good fire detection and suppression systems. The server farm contains four servers, each performing a specific data-processing function. The data is backed up weekly, but it takes time to restore, and there is a significant likelihood that some of the data will be lost due to the weekly backup schedule. Let's also say (just for the sake of illustration) that each of these servers costs $10,000 to replace with a new server (since they don't make your model any longer unless you were lucky enough to find a

cheaper one on eBay). You've attempted to quantify factors for the owners, such as *recovery time objective (RTO)* and *recovery point objective (RPO)*. These factors tell you how much data you could afford to lose if a disaster happened (such as a fire burning down that small wooden building) and how fast you would need to recover data to resume business operations to avoid suffering a significant loss of business.

Let's assume that the small business is in an area prone to thunderstorms, tornadoes, and hurricanes. One of the natural incidents that could occur in these areas is a lightning strike during a thunderstorm, where lightning strikes the metal roof of the building, and one of the servers suffers a power surge that completely renders it inoperable. If, after such an event, an assessment determines that the server can't be repaired, you decide to buy another one. The replacement cost of this server is $10,000. Let's do some simple quantitative analysis on this scenario:

- *Asset value (AV)* = $10,000
- *Exposure factor (EF)* = 1 (in this case, a 100% loss of the server and its data)
- *Single loss expectancy (SLE)* = Asset value (AV) × Exposure factor (EF) = $10,000

With an SLE of $10,000, you must hope this scenario never happens. However, based on an analysis of historical weather patterns in the area, and given the shoddy construction and limited electrical protection in the building where the server farm is located, this seems to be a likely possibility. So, you should determine how much the organization is likely to lose per year in such a scenario. You can calculate this if you know a couple of other pieces of information. First, you need to know how often severe storms are expected in the area (say, on an annual basis). This *annualized rate of occurrence (ARO)* could be used with the SLE value given earlier to determine your loss expectancy for a single year.

- SLE × ARO = *Annual loss expectancy (ALE)*

So, given that severe thunderstorms are historically shown to occur 15 times per year in your area (likelihood), you could conclude the following:

- $10,000 (SLE) × 15 (ARO) = $150,000 in potential loss per year (ALE)

Obviously, this is a bit of a stretch. Even if there are 15 severe lightning-producing storms per year, it's not likely that each one would produce a lightning strike on your server building every time. Even if that could happen, it's likely that after it happens the first time, your small business owners would invest in a new facility or relocate their servers, thus reducing the likelihood of other events. So, you might also do some statistical analysis to determine the likelihood of one or more of those 15 severe thunderstorms causing a lightning strike on the server farm building. This analysis may significantly reduce the likelihood of events from 15 to, say, 0.3–1 per year. Now your ALE would be a potential loss of only $3,333–$10,000 per year. Keep in mind that the numbers indicate the potential risk to the organization, which is expressed in terms of lost revenue and assets.

Again, the previous scenario is a simplistic (but not too unrealistic) example. Other factors could come into play that are not considered in the previous simple analysis. For example, suppose the business is losing $1,000 per day in revenue for every day this server is out of commission, which the previous calculations do not consider. This consideration is not only related to the SLE value (since it increases the SLE over time) but also influences your RTO and RPO values. You should also realize that this scenario considers only one of the servers; a better risk assessment should include the possibility of all assets being partially or completely lost.

Additionally, it's possible that a total loss of a server may not occur; let's say it's only a partial loss (say, 40%, or an EF of 0.4), which would reduce the SLE for that single asset ($10,000 [AV] × 0.4 [EF] = $4,000). So, you must take all of this into account when performing a quantitative risk analysis. There are many other factors that aren't mentioned here that could impact your SLE values (asset depreciation, lost labor hours to repair or replace the server, and so on). However, you can see the basics based on this simple example.

The final piece of this puzzle is determining the cost to the organization to reduce this risk. Let's examine all of the factors that the organization could influence. First, it could not affect the weather, so the threat agent and threat cannot be reduced or eliminated. However, the organization could address the vulnerability (the server building with its inadequate protection against bad weather, in this case), reducing it somewhat by relocating the server farm or even moving its applications and data to the cloud, for example. The organization could also reduce the likelihood of that particular threat-vulnerability pairing if it removed or changed the nature of the vulnerability itself, as described earlier.

The organization could also reduce the impact of this specific event in several ways. It could create a real-time failover cluster for the servers in another location that would instantly become available if said servers got struck by lightning, or it could at least have a hot spare server and a recent backup, if nothing else. The organization can mitigate the risk factors, but there is a cost for doing this. What is essential in developing risk responses is that the organization must balance the cost of mitigations (for example, acquiring another facility, moving to the cloud, using server clusters, or at least having hot spares and better backups) with how much it loses in terms of SLE and ALE and the added impact of lost revenue per day, as in this example. The bottom line is that you must find a solution that isn't more expensive or impacts the organization to a degree worse than the risk event itself. Risk responses are covered in more detail in Chapter 3.



Be familiar with the key elements of quantitative risk assessment techniques, such as asset value, exposure factor, single loss expectancy, annual rate of occurrence, and annualized loss expectancy.

Qualitative Risk Analysis

In a *qualitative assessment*, *qualitative* risk analysis involves using subjective scales. These scales would have subjective values (e.g., High, Medium, and Low). Qualitative risk analysis can be informed by historical trend analysis, experience, expert opinion, existing internal and external environmental factors, governance, and other inputs that are not always quantifiable but are nonetheless important. Qualitative risk analysis is appropriate when evaluating risk based on factors that are difficult to quantify, such as those not typically measured with concrete, repeatable values. For example, what value could you assign to the potential impact of the loss of consumer confidence or damage to reputation? You may attempt to frame it in terms of lost revenue (an educated guess at best) or loss of business. You could also survey a sample population to determine how many people would or would not do business with the organization after a data breach. However, even that would be subjective and offer limited insight into a hard-to-define measurement.

In Figure 2.7, you can see how the relationships of likelihood and impact affect different risk levels. The values of each are correlated, resulting in a qualitative risk determination. For example, according to this figure, a *moderate* level of likelihood combined with a *very high* level of impact results in a level of risk determined as *high*. Determining what constitutes a likelihood of moderate and what is required for an impact level of very high is subjective. However, it should be defined within the organization for consistency and standardization.

Semi-quantitative Risk Analysis

In *semi-quantitative* risk analysis, numeric scales are used to rate the probability and impact of risks.

Instead of values such as Low, Medium, and High, you might assign numerical values that roughly correspond to those levels. For example, you might use a type of *Likert scale*, which may have values such as 1 (Very Low), 2 (Low), 3 (Medium), 4 (High), and 5 (Very High). This might characterize the subjective level of impact or likelihood that the organization assigns to a particular risk event. In assigning these subjective values, you must also decide how they relate to each other, yielding an overall risk calculation. In other words, given likelihood values of 1 (Very Low), 2 (Low), 3 (Medium), 4 (High), and 5 (Very High), and identical values for impact, what would the overall risk level be when these two dimensions of risk are viewed together? How do their values affect each other?



NIST SP 800-30 refers to numerical scales as *semi-quantitative* values since they are numerical values that can be used in calculations but are still subjective in nature (i.e., they cannot be independently derived or repeated and are based on informed opinion).

Combining Quantitative and Qualitative Techniques

If you think that a qualitative or quantitative assessment alone may not be enough to adequately capture the different elements of risk and develop an overall risk determination, you're right. While there are instances where each of these techniques can be used alone,

they are typically used in combination to varying degrees in most situations. Each technique is good at producing a valuable piece of the puzzle, but sometimes you need both a quantitative perspective and a qualitative perspective on risk.

For example, in the scenario described earlier for quantitative risk, say you have numerical values for AV, EF, SLE, and ALE. These values account primarily for impact. However, suppose you don't have enough solid information to calculate likelihood. You might have to give it a "best guess" that comes from a reasonably well-thought-out analysis, using anecdotal data or another type of data that can't necessarily be quantified but is a reasonably accurate predictor. You may have to assign semi-quantitative likelihood values or the subjective High, Medium, and Low values. There might also be other subjective factors contributing to these values, so you would not have a purely quantitative assessment. Suppose you added "customer confidence" or "reputation" into the discussion on impact or loss to the company, for instance. Those would almost have to be qualitative values, so the assessment couldn't always be expressed in exact numerical terms that are easily repeatable or defendable.

In practice, most risk assessments are usually qualitative with some quantitative elements, or vice versa. There are usually few purely quantitative or entirely qualitative assessments performed. You should understand the elements and characteristics of both, however, for the exam and your profession.

Other Analysis Techniques

Now that you are familiar with the basics of qualitative and quantitative techniques, this section discusses a few specific techniques, based on both methods. Hence, you are well-prepared for the exam and success in a role in risk management. Note that professionals have developed dozens of different techniques over the years to assess risk; you'll find that most of them fall into either the qualitative or quantitative methodology. The ISO/IEC 31010:2019 standard, which was briefly described earlier, lists more than 30 such techniques in its appendices. Some of them approach assessments from different angles, but most share several commonalities. The following sections cover just a few of these.

Fault- and Event-tree Analysis

A *fault-tree analysis* is a method where the analyst starts with a risk event and looks for all possible causes. As the example in Figure 2.9 shows, this is diagrammed as a tree, with the risk event at the root and the causes extending away as branches. This analysis is used to discover all of the potential causes of the event (threats and threat actors in this case) to be further analyzed for likelihood.

Like fault-tree analysis, *event-tree analysis* is also diagrammed as a tree structure, but as seen in the example shown in Figure 2.10, it takes the opposite approach (bottom-up instead of top-down). In event-tree analysis, an initiating event is at the root of the tree, with all possible consequences (impacts) diagrammed from the root as branches (as opposed to causes of the event in fault-tree analysis). Both fault-tree and event-tree are predictive analysis tools; in other words, they seek to predict relationships between causes, events, and consequences.

FIGURE 2.9 An example of a fault-tree analysis diagram.

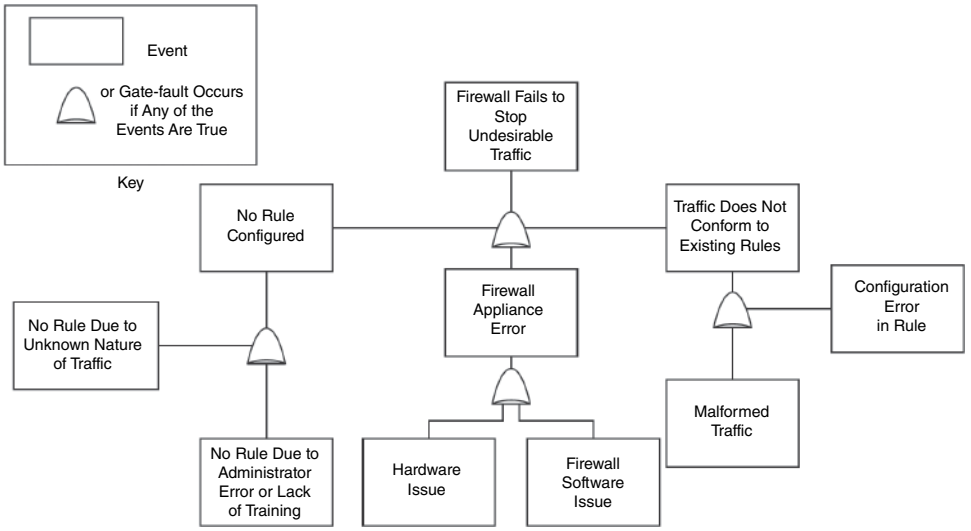


FIGURE 2.10 An example of an event-tree analysis diagram.

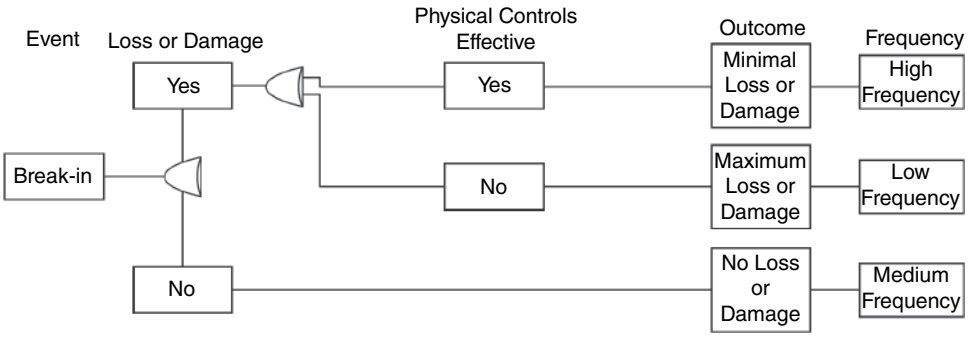


Figure 2.9 shows a risk event (in this case, a firewall failure), with the potential causes shown as no rule configured, a firewall appliance failure, and unknown or unusual traffic. Note that the risk event itself is the root of the tree. Conversely, in Figure 2.10, an event-tree diagram, the risk event (in this case, a break-in to a corporate facility) is also at the root. However, the branches leading away are the potential consequences, not the causes.

Factor Analysis of Information Risk (FAIR)

The *Factor Analysis of Information Risk (FAIR)* is a semi-quantitative methodology used to determine the probable loss event frequency and magnitude for individual risks.

The FAIR methodology consists of four stages:

- **Identify scenario components:** This stage identifies the asset at risk, as well as the threat community being analyzed.

- **Evaluate loss event frequency:** Activities in this stage include an estimation of threat event frequency, threat capability, protective control strength, vulnerabilities, and the frequency of loss events.
- **Evaluate probable loss magnitude:** This stage is used to develop worst-case and most likely loss magnitudes.
- **Derive and articulate risk:** Derived from the values established in earlier stages, the level of risk is calculated and portrayed.

More facts about FAIR include:

- FAIR provides a quantitative risk analysis methodology that can be integrated into the Assess and Monitor steps in NIST 800-37.
- NIST SP 800-30 defines qualitative and quantitative assessment concepts but does not prescribe a method; FAIR fills that gap by providing a structured model for quantifying loss event frequency and magnitude.
- Organizations can use FAIR outputs (loss frequency, loss magnitude, probable risk) to express risk in financial terms that align with RMF's risk determination and acceptance steps.
- FAIR is published by the Open Group as the Open FAIR™ Risk Analysis (O-RA) and Risk Taxonomy (O-RT) standards.
- These standards define the taxonomy and model for consistent information risk measurement across organizations.
- The Open FAIR framework is often used in conjunction with NIST, ISO/IEC 27005, and ISACA's Risk IT, serving as a quantitative extension to these frameworks rather than a replacement for them.

The FAIR methodology is presented in the book *Measuring and Managing Information Risk: A FAIR Approach*, by Jack Freund and Jack Jones. The FAIR Institute, an organization offering membership, training, and events, can be found at `www.fairinstitute.org`.

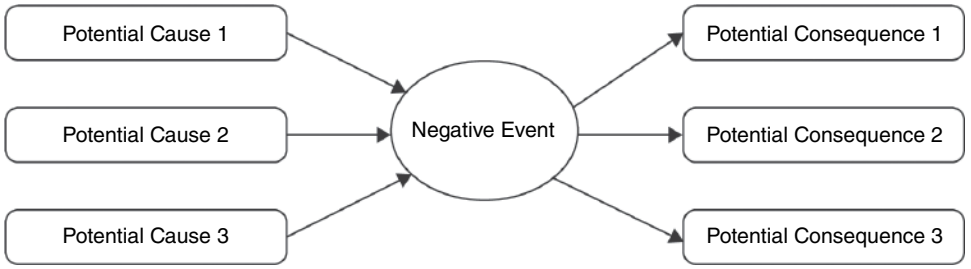
Bow-tie Analysis

A *bow-tie analysis* uses diagrams to analyze and explain relationships between various risk elements, from causes (threats) to events and then to impacts (consequences). It is similar to both the fault-tree analysis and the event-tree analysis. It looks at the various causes of a risk event (fault tree) and analyzes the consequences of the event (event tree). The difference, however, is that the bow-tie analysis looks at the intervening characteristics of the events and causes, such as the path by which the cause leads to the event, and then the consequences. Figure 2.11 illustrates a bow-tie diagram. In this figure, the adverse event is shown as the center of the bow tie, with potential causes on the left and possible consequences on the right.

Other Methods

As mentioned, there are numerous other methods for performing the different types of assessments, which cannot be discussed here. Some of these methods focus on data collection, while others emphasize different group analysis techniques. All are, again, some

FIGURE 2.11 An example of a bow-tie analysis diagram.



variation of either the qualitative or quantitative method, and may require examining specific aspects of data to determine various elements of risk. Other examples of these techniques include the Delphi method, which centers on expert opinion from questionnaires; the Bayesian analysis method, which uses data distribution and statistical inference to determine risk probability values, checklists, scenario analysis, and business impact analysis (described later in this chapter); root cause analysis (discussed earlier in this chapter); and various collaborative techniques, such as *brainstorming*. You won't be expected to know all these methods for the exam, but you should be familiar with the concepts of qualitative and quantitative analysis and some of the techniques that each uses. For further knowledge on various assessment and analysis techniques, consult with IEC 31010:2019, available from `iso.org`.



You will not need to know any particular risk analysis method in detail for the exam, but you should understand how the qualitative, quantitative, and semi-quantitative methods work.

Business Impact Analysis

A *business impact analysis (BIA)* is used to identify an organization's business processes, the interdependencies between processes, the resources required for process operation, and the impact on the organization if any business process is incapacitated for a time for any reason.

While a BIA is considered a cornerstone of a *business continuity* and *disaster recovery* (BCDR) program, it is also useful for risk analysis. The BIA reveals the most critical processes (and underlying resources, such as information systems and other IT infrastructure), giving the risk manager a better understanding of which business processes and systems warrant the most protection from an operational criticality perspective.

The presence of a BIA provides a strong indication of the organization's maturity through its intention to protect its most critical processes from disaster scenarios. Correspondingly, the absence of a BIA may suggest that the organization does not consider risk management as having strategic importance.

Inventory Key Processes and Systems

The first step in a BIA is the collection of key business processes and IT systems. Within the overall scope of the business continuity planning (BCP) project, the objective here is to establish a detailed list of all identifiable business processes and systems. The usual approach is to develop a questionnaire or intake form that would be circulated to key personnel in end-user departments and IT. Figure 2.12 shows a sample intake form.

Typically, the information collected on intake forms is transferred to a multi-columned spreadsheet, where information on all of the organization's in-scope processes can be viewed together. This will become even more useful in subsequent phases of the BCP project, such as the criticality analysis.



The use of an intake form is not the only accepted approach when gathering information about critical processes and systems. It's also acceptable to conduct one-on-one or group interviews with key users and IT personnel to identify critical business processes and systems. I recommend using an intake form (whether paper-based or electronic), even if the interviewer uses it as a framework for notetaking.

FIGURE 2.12 BIA sample intake form for gathering data about key processes.

Process or system name	
Interviewee	
Title	
Department	
Contact info	
Date	
Process owner	
Process operator(s)	
Process description	
Customer facing (Y or N)	
IT system(s) used	
Key suppliers	
Communications needed	
Assets needed	
Process dependencies	
Other dependencies	
Documentation location	
Records location	

Planning Precedes Action

IT personnel in retail organizations are often eager to get to the fun and meaty part of a project. Developers are anxious to begin coding before design is completed, reviewed, and approved. System administrators are eager to build out the infrastructure before it is scoped and designed. BCP personnel fervently desire to start designing more robust system architectures and tinker with replication and backup capabilities before architecture and design are completed.

In business continuity and disaster recovery planning, completing the BIA and other analyses is critical, as the analyses help define the systems and processes most needed before getting to the fun part.

Governance structures must define and enforce the sequence of events (business case, requirements, architecture, design, development, etc.) to ensure favorable outcomes.

Statements of Impact

When processes and systems are being inventoried and cataloged, it is also vitally important to obtain one or more statements of impact for each process and system. A *statement of impact* is a qualitative or quantitative description of the impact on the business if the process or system were incapacitated for a time.

For IT systems, you might capture the number of users and the names of departments or functions affected by the unavailability of a specific IT system. Include the geography of affected users and functions if that is appropriate. Here are some sample statements of impact for IT systems:

- Three thousand users in France and Italy will be unable to access customer records, resulting in degraded customer service.
- All users in North America will be unable to read or send email, resulting in productivity slowdowns.

Statements of impact for business processes might cite the business functions that would be affected. Here are some sample statements of impact:

- Accounts payable and accounts receivable functions will be unavailable for processing, impacting the availability of services and supplies and reducing revenue.
- The legal department will be unable to access contracts and addenda, resulting in lost or delayed revenue.

Statements of impact for revenue-generating and revenue-supporting business functions could quantify financial impact per unit of time (be sure to use the same units of time for all functions to be easily compared with one another). Here are some examples:

- The inability to place orders for materials will incur a cost of $12,000 per hour.
- Delays in payments will cost $45,000 per day in interest charges.

As statements of impact are gathered, it might make sense to create several columns in the main worksheet so that like units (names of functions, numbers of users, financial figures) can be sorted and ranked later.

When the BIA is completed, you'll have the following information about each process and system:

- Name of the process or system
- Who is responsible for its operation
- A description of its function
- Dependencies on systems
- Dependencies on suppliers
- Dependencies on key employees
- Quantified statements of impact in terms of revenue, users affected, and/or functions impacted

Criticality Analysis

When all the BIA information has been collected and charted, the criticality analysis (CA) can be performed. *Criticality analysis* is a study of each system and process, considering the impact on the organization if it is incapacitated, the likelihood of incapacitation, and the estimated cost of mitigating the risk or impact of incapacitation. In other words, it's a somewhat special type of risk analysis that focuses on key processes and systems.

The criticality analysis needs to include, or reference, a threat analysis. A *threat analysis* is a risk analysis that identifies every threat with a reasonable probability of occurrence, along with one or more mitigating controls or compensating controls, and the new probabilities of occurrence with those mitigating/compensating controls in place. To give you an idea of what this looks like, refer to Table 2.6, which provides a lightweight example.

In the threat analysis presented in Table 2.6, notice the following:

- Multiple threats are listed for a single asset. The preceding example mentioned just eight threats. For all the threats but one, only a single mitigating control is listed. For the extended power outage threat, two mitigating controls are listed.
- The cost of downtime wasn't listed. For systems or processes that incur a cost per unit of time for downtime, you'll need to include it here, along with calculations to demonstrate the payback for each control.
- Some mitigating controls can benefit more than one system. That may not have been obvious in this example. However, in the case of an uninterruptible power supply (UPS) and an electric generator, many systems can benefit, so the cost for these mitigating controls can be allocated across many systems, thereby lowering the cost for each system. Another example is a high-availability storage area network (SAN) located in two different geographic areas. Although it's initially expensive, many applications can use the SAN for storage, and all will benefit from replication to the counterpart storage system.

TABLE 2.6 Sample threat analysis for identifying threats and controls for critical systems and processes.

System: Application Server				
Threat	**Probability**	**Mitigating Control**	**Mitigation Cost**	**Mitigated Probability**
Denial of service	0.1%	High-performance filtering router	$60,000	0.01%
Malware	1%	Antivirus	$200	0.1%
Storage failure	2%	RAID 5	$20,000	0.01%
Administrator error	15%	Configuration management tools	$10,000	1%
Hardware CPU failure	5%	Server cluster	$15,000	1%
Application software bug	5%	Source code reviews	$10,000	2%
Extended power outage	25%	UPS Electric generator	$12,000 $40,000	2% 0.5%
Flood	2%	Relocate data center	$200,000	0.1%

- Threat probabilities are arbitrary. In Table 2.6, the probabilities are for a single occurrence in an entire year (for example, 5% means the threat will be realized once every 20 years).
- The length of the outage was not included. You may need to include this, particularly if you are quantifying downtime per hour or another unit of time.

It is probably becoming evident that a threat analysis and the corresponding criticality analysis can get complicated. The rule here should be as follows: the complexity of the threat and criticality analyses should be proportional to the value of the assets (or revenue, or both). For example, in a company where application downtime is measured in thousands of dollars per minute, it's probably worth taking a few weeks or even months to work out all the likely scenarios and a variety of mitigating controls, and then work out which ones are the most cost-effective. On the other hand, for a system or business process where the impact of an outage is far less costly, a good deal less time might be spent on the supporting threat and criticality analysis.

Recovery Objectives and Next Steps

After the criticality analysis is performed, business continuity planners will establish a metric known as *maximum tolerable downtime (MTD)*. This is the amount of time that the organization can tolerate a business process being nonfunctional before the ongoing viability of the business itself (or a major part of it) is at risk. Note that MTD is a theoretical figure that the organization will not attempt to verify.

Business impact analysis is often used to establish recovery objectives that drive the development of disaster recovery plans. Recovery objectives and disaster recovery planning are discussed in more detail in Chapter 4.

Inherent and Residual Risk

Coverage of the IT Risk Assessment domain in the CRISC certification continues with a discussion of inherent risk and residual risk. These two concepts introduce a small degree of added complication to the overall risk management and risk treatment concepts.

Risk treatment is discussed in full detail in Chapter 3.

Inherent Risk

No personal or business activity is entirely risk-free, and some activities are inherently riskier than others. The concept of *inherent risk* refers to the level of risk associated with a process or activity before applying any protections or safeguards. In the vernacular, some activities are just riskier than others.

Take, for example, two personal activities: gardening and hang gliding. I would state that the inherent risk associated with hang gliding is significantly higher than that of gardening. With hang gliding, any number of realistic scenarios can result in the death of the hang glider. At the same time, gardening is relatively safe, and death by gardening is highly unlikely. In both gardening and hang gliding, the probability of risk realization is relatively low, although the impact of risk realization is decidedly higher for hang gliding.

One approach to risk analysis is assigning inherent risk to a process or system, and then calculating changes to risk after protective controls and safeguards are applied. This helps a risk analyst understand the changes in risk when individual controls and safeguards are applied, and it better informs management of the full nature of risk for a given part of the organization.

Residual Risk

The risk management and risk treatment processes tend to reduce risk, but they rarely eliminate all risk. Instead, risk mitigation and risk transfer reduce risk to a lower level. This "leftover" risk is known as *residual risk*.

Residual risk is best explained with a real-world example. An organization's workforce is equipped with laptop computers running Windows or macOS, neither of which utilizes whole-disk encryption. Risk analysis suggests that the risks associated with the theft or disappearance of an unencrypted laptop computer are quite high, with potential for a significant security or privacy incident, depending on whose laptop computer is lost or stolen.

In this scenario, the organization decides to implement whole-disk encryption on all its laptop computers. Note that this measure does nothing to change the probability of laptop computers being lost or stolen. However, this measure significantly alters the impact. While whole-disk encryption ensures the organization no longer needs to worry about potential data compromise, the stolen laptop will still need to be replaced. The residual risk, in this example, refers to the financial risk associated with the funds required to purchase a new laptop computer, as well as any labor needed to prepare it for employee use.

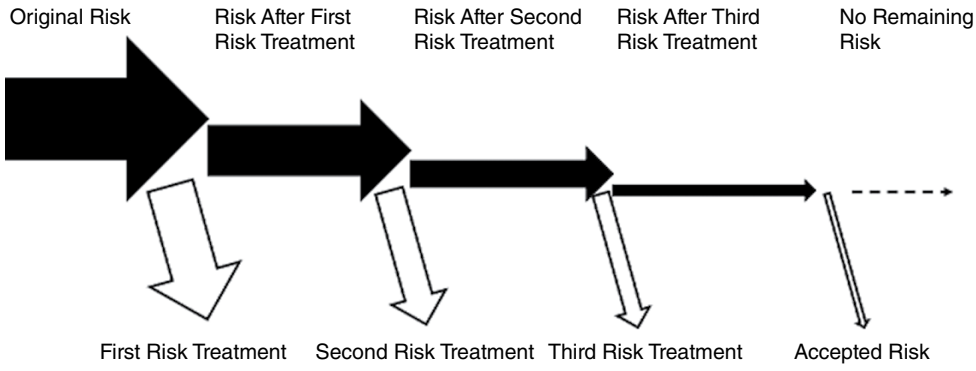
To continue this example, the residual risk may still be a risk that warrants further analysis and treatment. Analysis of the risks associated with the need to replace a lost or stolen laptop computer could result in a recommendation of the purchase of security cables that are issued to each employee, together with directives that laptop computers are to be locked to a piece of furniture (or locked in a safe) when not in the direct control of the employee. This measure reduces the probability of theft or loss, but not the impact: if the laptop computer is stolen, it must still be replaced, but theft is less likely to occur since, at least some of the time, employees will use those cable locks, reducing the number of opportunities for someone to steal them.

Larger and more complex risk scenarios may require multiple rounds of risk analysis and risk treatment. However, an organization will eventually reach the point where the residual risk is low enough that business leaders and department heads will accept it.

Often, risk analysis, risk treatment, and residual risk are more complicated than this. A risk assessment performed on a business process or information system is likely to identify multiple risks, which should be considered both collectively and individually. It may be easier to compare various types of risks and risk treatment options when they are considered together. Sometimes, a single risk treatment measure may mitigate more than one of the identified risks. Figure 2.13 depicts multiple rounds of risk analysis and risk treatment until the final remaining risk is accepted.

Let's go back to the stolen laptop example for a moment. Imagine that multiple risks are identified in a stolen laptop scenario, all having to do with the compromise of

FIGURE 2.13 Risk reduction through multiple rounds of risk treatment.



different types of information: *personally identifiable information (PII)* such as human resource information, non-public financial information, intellectual property, and trade secrets. Whole-disk encryption will reduce the impact of all four of these scenarios to near zero.

Miscellaneous Risk Considerations

There may be other options to consider when performing risk assessments, which may or may not be applicable in an organization. Understanding industry and compliance models related to an organization's industry can be absolutely critical to using the risk assessment to pass an audit or ensure legal compliance. There may be unique methodologies used by the organization to illustrate different aspects of a cybersecurity program.

For companies that handle *protected health information (PHI)*, there are additional legal considerations related to the data, as these companies must comply with the *Health Insurance Portability and Accountability Act (HIPAA)* regulations. In this case, HIPAA requires tying each risk to confidentiality, integrity, and availability. This can be as simple as creating a new column on a risk register that denotes if confidentiality, integrity, or availability is affected by the risk. It is a simple additional step that can help demonstrate not only compliance with HIPAA, but a deeper understanding of how the risks may affect an organization.

The enforcement arm of HIPAA is the Office of Civil Rights (OCR). Working with their regulations, they have risk requirements that need to be assessed for legal compliance reasons. The OCR does not expect companies to follow their exact risk methodology, as long as the concerns are fully applied within the organization's risk framework. If there ever is a HIPAA violation where the OCR becomes involved, considerations like this and assessing confidentiality, integrity, and availability can make regulators feel more at ease related to the violation and sometimes help reduce the overall penalties.

Stepping away from HIPAA considerations, sometimes organizations may choose to follow or report to others using frameworks such as NIST CSF (described previously). In this case, as with the confidentiality, integrity, and availability example, the NIST CSF functions (Identify, Protect, Detect, Respond, Recover, and Govern) may be applied similarly. This can sometimes be useful for helping business executives who do not understand the intricacies of cybersecurity to grasp, at least at a high level, why certain risks are so important to consider.

Another option that may or may not be helpful, depending on the organization, is the concept of *materiality*. Materiality originates from the audit world and refers to a significant or fundamental flaw in a specific program. For example, if an organization lacks antivirus software, that is a fundamental flaw and thus a material weakness. Adding a column like this to a risk register is yet another way that organizations can create extra emphasis around a risk management program.

It should be clear that there are several ways organizations can approach risk. In some instances, it is not so much about being right or wrong concerning a risk methodology but rather about working through the right flavor of risk within an organization.

Using AI to Identify and Analyze Risk

Organizations are using *artificial intelligence (AI)* and *machine learning (ML)* to transform how they identify, assess, and monitor risk. Traditional risk assessments rely on human capabilities and judgment, periodic data gathering, and manual analysis—approaches that can overlook rapidly changing threat conditions. AI can continuously collect, correlate, and interpret signals from internal and external data sources to detect patterns, anomalies, and emerging risks at a scale and speed that exceeds human capacity. When used properly, AI can enhance the quality and responsiveness of IT risk management without replacing human expertise or governance oversight.

Role of AI in Risk Identification

AI tools can automatically identify indicators of potential risk in both structured and unstructured data. In risk identification, AI systems utilize pattern recognition, natural language processing (NLP), and predictive analytics to analyze large volumes of operational, network, and environmental data, thereby revealing new risk scenarios. Examples include:

- **Threat intelligence integration:** AI algorithms can aggregate and classify information from threat feeds, vulnerability databases, and open-source intelligence feeds to identify relevant emerging threats.
- **Anomaly detection:** ML models learn normal behavior within systems and processes, then flag deviations, such as unusual login patterns, abnormal data transfers, or process delays that may indicate risk events.
- **Text and sentiment analysis:** NLP engines can analyze incident reports, audit findings, or even social media and supplier communications to detect indicators of operational, compliance, or reputational risk.
- **Asset discovery and mapping:** AI-driven tools can continuously inventory and classify digital assets, including cloud resources and endpoints, thereby reducing the likelihood of unknown or "shadow" assets that introduce dark risk.

These examples align with NIST SP 800-30, which specifies that risk identification should include the discovery of threat sources, events, vulnerabilities, and predisposing conditions. AI enhances some of these tasks by automating data collection and correlation, enabling risk professionals to identify threats and vulnerabilities that would otherwise remain undetected.

AI in Risk Analysis and Evaluation

Once risks have been identified, AI tools can support quantitative and qualitative analysis through probabilistic modeling and data-driven estimation. Within the NIST Risk Management Framework (SP 800-37), this activity corresponds to the *Assess* and *Monitor* steps. In ISO/IEC 27005:2022 and IEC 31010:2019, it relates to risk analysis and evaluation.

Key AI analysis and evaluation applications include:

- **Predictive modeling:** ML algorithms trained on historical incident and loss data estimate likelihood and impact more consistently than subjective expert judgment alone, enabling a more reliable risk triage.
- **Causal and correlation analysis:** AI systems can identify relationships among multiple risk factors. For example, linking supplier delays, patch latency, and service availability to uncover compound or systemic risks.
- **Control effectiveness evaluation:** Automated analytics can test control performance using continuous monitoring data, which reveals deviations that would not be detected in periodic audits.
- **Quantitative estimation:** When paired with quantitative methodologies such as the Factor Analysis of Information Risk (FAIR) model, AI can automate calculations of probable loss frequency and magnitude, thereby improving repeatability and accuracy.

AI systems can also maintain and update a dynamic risk register. By continuously analyzing operational data, they can detect changes in exposure, control performance, or threat environment, and automatically adjust risk ratings. This supports the *risk monitoring* principle in both NIST and ISO standards, turning the static risk register into a living, data-driven record of an organization's risk posture.

Integration with Frameworks and Standards

AI does not replace established risk management frameworks, but extends their analytical power. Examples include:

- **NIST RMF:** AI enhances ongoing assessment, continuous monitoring, and reporting. Automated analysis of security control data supports NIST SP 800-53A control assessments.
- **ISO/IEC 31000:2018:** AI contributes to evidence-based risk evaluation and monitoring by generating data-driven insights while maintaining the framework's emphasis on context, communication, and governance.
- **ISACA Risk IT and COBIT:** AI capabilities can integrate with Risk IT's Collect Data and Analyze Risk processes (RE1 and RE2), providing real-time analytics to support decision-making.
- **FAIR (Open Group Standard):** AI techniques can enhance the FAIR quantitative approach by improving parameter estimation and data normalization for loss event frequency and magnitude.

Benefits and Limitations

Some of the benefits of using AI in risk assessments and analysis include:

- Faster and more comprehensive identification of emerging threats and vulnerabilities
- Improved consistency and repeatability of risk scoring and prioritization

- Real-time monitoring of control effectiveness and risk exposure
- Enhanced visualization and communication of complex risk relationships

Some of the limitations and risks in the use of AI for risk assessments and analysis include:

- **Dependence on automation:** Many of the AI-related capabilities listed in this section depend on automation, including the use of centralized event logging, processing, and response, configuration management data, and threat intelligence feeds. Before AI can enhance these and other capabilities, it must first exist with some degree of automation.
- **Model bias and opacity:** AI systems can reflect bias present in training data or algorithms, leading to skewed risk assessments and results.
- **Data quality dependence:** Inaccurate, incomplete, or unbalanced datasets degrade model accuracy, skewing results.
- **Explainability and accountability:** Decisions derived from opaque models may conflict with governance and audit requirements, unless supplemented with transparent rationale.
- **Ethical and privacy considerations:** AI-driven data aggregation and inference may raise privacy or fairness concerns, requiring alignment with organizational ethics policies and applicable laws.

Governance and Human Oversight

AI should augment, not replace, professional judgment in risk management. Effective use of AI requires strong governance controls to validate algorithms, monitor performance, and maintain accountability. Risk managers must ensure that AI tools comply with internal governance standards, including model validation, data provenance, data protection, and change management. Results produced by AI systems should be reviewed by qualified analysts before influencing strategic or operational decisions through human-in-the-loop safeguards.

AI and machine learning introduce transformative capabilities for identifying, analyzing, and monitoring IT risk. By embedding AI within established frameworks, organizations can move from periodic, manual assessments to continuous, data-driven risk intelligence. However, AI models must be governed rigorously to ensure accuracy, fairness, and transparency. When properly managed, AI enhances the precision, agility, and business relevance of risk assessment, enabling organizations to anticipate and mitigate threats more effectively.

Summary

IT risk identification involves activities that aim to identify various types of risks associated with an organization's use of information technology, including strategic, operational, cybersecurity, privacy, supply chain, and compliance risks. The types of activities where risks may be identified include risk assessments, audits, incidents, penetration tests, news and social media, advisories, threat modeling, word of mouth, whistleblowers, and passive observation.

Threat modeling is a risk identification technique that involves examining every possible threat agent, action, event, attack vector, and vulnerability for a given system, asset, or process, and then modeling or simulating how it could progress and the damage that could occur. Vulnerability and control deficiency analysis are crucial components of risk analysis, focusing on weaknesses in a system or process that could lead to a potentially serious incident if exploited by a threat actor. *Root cause analysis*, or RCA, is a method of problem-solving that seeks to identify the root cause of an event, situation, or problem.

A *risk scenario* is a potential real-world, business-impacting event that consists of a threat, threat actor, vulnerability, and asset. *Risk analysis* is a detailed examination of a risk to better understand its basic factors, including the likelihood and impact of risk occurrence, and the development of a measure that might be implemented to reduce the likelihood or impact.

Several types of assessments can be used to identify risks, including a gap assessment, threat modeling, vulnerability assessment, maturity assessment, penetration test, data discovery, code review, architecture review, design review, code scan, audit, and risk assessment. Several techniques are used in a risk assessment to identify risks, including interviews, documentation review, observation, and system testing.

Risk ranking is the process of sequencing risks, generally with the greatest risks appearing first. A *risk owner* is the individual responsible for a business process or system that focuses on mitigating an identified risk. A *risk register* is a business record containing information about business risks and their origin, potential impact, affected assets, probability of occurrence, and treatment.

A *risk analysis* is the detailed analysis of a risk that has been identified in an activity, such as a risk assessment. *Quantitative risk analysis* uses concrete (nonsubjective) numerical values and statistical analysis to calculate the likelihood and impact of risk. On the other hand, a *qualitative risk analysis* involves using subjective scales, such as a numeric scale from 1 to 10, or subjective values, such as High, Medium, and Low. There are several risk analysis techniques, including fault- and event-tree analysis, Factor Analysis of Information Risk (FAIR), and bow-tie analysis.

A *business impact analysis (BIA)* is used to identify an organization's business processes, the interdependencies between processes, the resources required for process operation, and the impact on the organization if any business process is incapacitated for a time for any reason. A *statement of impact* is a qualitative or quantitative description of the impact on the business if the process or system were incapacitated for a time. A *criticality analysis* is a study of each system and process, considering the impact on the organization if it is incapacitated, the likelihood of incapacitation, and the estimated cost of mitigating the risk or impact of incapacitation.

Inherent risk refers to the level of risk associated with a specific activity or function before any protective controls or safeguards are applied. *Residual risk* is the risk that remains after any risk transfer or risk mitigation measures have been implemented. The *materiality* of a risk may help to highlight its importance to management.

AI may be used in risk identification and risk analysis, provided that proper AI governance safeguards are in place. While recognized as a force multiplier in business and industry, imprudent uses of AI may result in regulatory scrutiny and reputational damage.

Exam Essentials

Understand the types of IT risks and the ways they may be identified. The types of IT risks include strategic, operational, cybersecurity, privacy, supply chain, and compliance. Risks may be identified through various means, including audits, incidents, news and social media, security advisories, whistleblowers, and risk assessments.

Understand the unique vocabulary of risk and risk assessments. Key terms in risk assessment and analysis include asset, vulnerability, threat, impact, attack, and risk. Additional terms include threat actor, threat event, threat modeling, exploit, inherent risk, residual risk, probability, risk register, risk analysis, risk assessment, risk treatment, and risk management.

Understand the purpose of vulnerability and control deficiency analysis. Vulnerability analysis is an examination of systems or processes to identify vulnerabilities that a threat agent could exploit to harm an asset. Control deficiency analysis is an evaluation of a control to determine whether it is effective and meets its intended control objective.

Understand the purpose and technique of root cause analysis. Root cause analysis (RCA) is a specific technique used to identify the ultimate cause of an event or incident. The core of RCA is the "five whys," in which staff members start with the event and work backward to arrive at the cause.

Understand the various types of assessments and their intended purposes. There are various assessments, including risk assessments, gap assessments, threat modeling, vulnerability assessments, maturity assessments, penetration tests, data discovery, design reviews, code reviews, code scans, audits, and more.

Understand industry standard risk assessment and risk analysis frameworks. Standards and frameworks include NIST SP 800-30 (risk assessments), OCTAVE (Operationally Critical Threat, Asset, and Vulnerability Evaluation), OCTAVE-S, OCTAVE-Allegro, ISO/IEC 27005, IEC 31010, and Risk-IT.

Understand the purpose and use of a risk register. A risk register is a business record that stores various characteristics of an identified risk. A risk's characteristics lend themselves to categorization, allowing for a better understanding of the big picture of risk across the organization.

Understand risk analysis methodologies. Common risk analysis methodologies include qualitative risk analysis, quantitative risk analysis, and semi-quantitative risk analysis.

Understand the role of business impact analysis (BIA), statements of impact, and criticality analysis. Business impact analysis is used to understand the impact on an organization if a system or process were to be unavailable. Statements of impact describe, in business terms, the consequences of a system or process being unavailable. A criticality analysis further defines the impact of an unavailable system or process.

Understand the concepts of inherent risk and residual risk. Inherent risk refers to the level of risk associated with a given process or activity before any safeguards are applied. Residual risk is the leftover risk after an identified risk has been treated (risk treatment rarely eliminates risk entirely).

Review Questions

The following questions are designed to test your understanding of this chapter's material. For more information on how to obtain additional questions, see this book's Introduction. You can find the answers at the end of this chapter.

1. An organization experiences a brief system outage that does not exceed established impact thresholds and requires no escalation. How should this occurrence be classified for risk management purposes?
   - **A.** A near miss with no relevance to risk analysis
   - **B.** A vulnerability discovery
   - **C.** A low-impact risk event
   - **D.** A control deficiency
2. Which example best illustrates a risk event rather than a threat or vulnerability?
   - **A.** Phishing emails targeting employees
   - **B.** Lack of multifactor authentication on remote access
   - **C.** Unauthorized access to customer data through compromised credentials
   - **D.** Absence of a formal access control policy
3. What is the maximum period of time that risk items should reside in a risk register?
   - **A.** One year
   - **B.** Seven years
   - **C.** Until the risk is accepted
   - **D.** Indefinitely
4. A risk manager reviewing a business process has identified a control gap. What steps should the risk manager take to determine the best course of action?
   - **A.** Perform a risk analysis to determine the magnitude of the risk
   - **B.** Determine whether the control gap represents a significant risk
   - **C.** See whether a risk already exists in the risk register
   - **D.** Create an entry in the security incident log
5. A security professional is performing a test with fuzzers and protocol analyzers. What is the most likely target of this test?
   - **A.** Application vulnerability analysis
   - **B.** Firewall rules review
   - **C.** Data discovery
   - **D.** Threat modeling

6. What is the primary purpose of risk scenario development?
   - **A.** To identify the highest-rated risks
   - **B.** To quantify the highest-rated risks
   - **C.** To prepare an executive briefing
   - **D.** To associate risks with business objectives

7. A risk manager seeks to gain a deeper understanding of the specific factors within particular business processes that could lead to an incident. What artifact would inform the risk manager of the urgency to perform this analysis?
   - **A.** Risk appetite
   - **B.** Privacy policy
   - **C.** Business impact analysis
   - **D.** Risk register

8. Hackers, cybercriminal organizations, disgruntled employees, and hurricanes are known as what?
   - **A.** Threat agents
   - **B.** Threat events
   - **C.** Threats
   - **D.** Attacks

9. A risk manager is considering using an industry standard for performing risk assessments. Which of the following should the risk manager NOT consider?
   - **A.** NIST RMF
   - **B.** NIST SP 800-30
   - **C.** OCTAVE
   - **D.** IEC 31010

10. What does NIST Special Publication 800-30 describe?
    - **A.** A risk management methodology
    - **B.** A risk assessment methodology
    - **C.** Security and privacy controls
    - **D.** Audit procedures for security and privacy controls

11. Why is it important to assess vulnerabilities and control deficiencies together?
    - **A.** To eliminate the need for threat modeling
    - **B.** To ensure all vulnerabilities are immediately remediated
    - **C.** To understand how weaknesses affect risk exposure in context
    - **D.** To prioritize technical controls over administrative controls

12. Which of the following is the best description of a risk register?

- **A.** A list of threat agents and threat actors
- **B.** A list of identified risks
- **C.** The root cause analysis (RCA) for a security incident
- **D.** A list of likely risk scenarios

13. A risk manager performs risk analysis on a specific risk to better understand it and identify potential remedies. Various aspects of the initial risk and its potential remedies are scored on a scale of 1–10. What type of risk analysis is being performed?

- **A.** Qualitative risk analysis
- **B.** Quantitative risk analysis
- **C.** Semi-quantitative risk analysis
- **D.** Factor Analysis of Information Risk (FAIR)

14. An organization has performed a risk analysis of a complex scenario. Two different mitigation strategies have been approved that will reduce the original risk to a much lower level. The risk that remains after these mitigation strategies have been completed is known as what?

- **A.** Inherent risk
- **B.** Leftover risk
- **C.** Residual risk
- **D.** Accepted risk

15. Which of the following is the best method for managing residual risk?

- **A.** Subject it to risk analysis and risk treatment
- **B.** Retain it on the risk register and consider it as accepted
- **C.** Retain it on the risk register and label it as closed
- **D.** Add it to the cumulative total of risk tolerance

16. Which activity best supports risk identification rather than risk analysis?

- **A.** Estimating loss expectancy using quantitative data
- **B.** Evaluating control effectiveness
- **C.** Discovering potential threat actors and vulnerabilities
- **D.** Assigning risk scores in a risk matrix

17. In a semi-quantitative risk analysis, a risk manager assigns 4 = Likely for likelihood and 5 = Severe for impact, producing a risk score of 20. What is the primary advantage of this approach?

- **A.** It eliminates the need for expert judgment
- **B.** It allows consistent ranking of risks without requiring precise financial data
- **C.** It converts qualitative data directly into monetary loss
- **D.** It ensures that results are statistically valid

18. What is the main purpose of a business impact analysis (BIA) within the IT risk assessment process?
    - **A.** To calculate quantitative loss expectancy
    - **B.** To identify critical business processes and the consequences of their disruption
    - **C.** To assign risk ownership across business units
    - **D.** To verify the effectiveness of existing controls
19. Which change most directly indicates a shift in an organization's threat landscape?
    - **A.** Implementation of a new security awareness program
    - **B.** Emergence of new attacker techniques targeting the organization's industry
    - **C.** Completion of an internal control self-assessment
    - **D.** Adoption of a new enterprise risk management framework
20. Which of the following best describes the purpose of threat modeling in an IT risk assessment?
    - **A.** To simulate the financial impact of business disruptions
    - **B.** To map potential attack paths and identify exploitable vulnerabilities in systems or processes
    - **C.** To test system controls through active exploitation
    - **D.** To document business impact and recovery objectives

Answers to Review Questions

1. C. Even when impact thresholds are not exceeded, the occurrence of an outage constitutes a realized risk event, albeit with low impact, and should be considered in ongoing risk analysis and trend evaluation.
2. C. Unauthorized access to customer data represents the realized outcome of a threat exploiting a vulnerability, making it a risk event rather than a threat condition or control weakness. The other choices are vulnerabilities or threats.
3. D. Individual risk register items should remain in the risk register for as long as the risk they describe continues to exist—indefinitely, even.
4. A. When a control gap is identified, the best first step is to perform a risk analysis to determine what risks (including compliance risk) exist that warrant consideration and action.
5. A. Fuzzers are most often used in an application vulnerability analysis. The role of a fuzzer is to present numerous inputs to input fields to determine whether the application is vulnerable to malicious inputs.
6. D. Risk scenario development is performed to associate specific risks with specific business objectives or business activities. Risk scenario development makes risks more tangible by describing their potential impact on the organization in business terms.

7. C. A business impact analysis (BIA) is used to determine the relative criticality of each business process in the organization. If the BIA determines that any particular business process is critical, this will signal that the business process warrants additional attention.

8. A. Threat agents are the entities, human or not, that may perform actions that could cause harm to an organization, its processes, or its systems.

9. A. NIST RMF is a framework for overall risk management, but not for risk assessments. NIST SP 800-30, OCTAVE, and IEC 31010 are all risk assessment standards that warrant consideration.

10. B. NIST SP 800-30, "Guide for Conducting Risk Assessments," is a step-by-step guide to performing qualitative and quantitative risk assessments. The standard also briefly describes the overall risk management lifecycle process to provide context for individual risk assessments.

11. C. Assessing vulnerabilities and control deficiencies together provides context for how weaknesses interact with existing controls and influence overall risk exposure.

12. B. A risk register is a master list of all credible risks that have been identified and are awaiting further analysis, treatment, or follow-up.

13. C. Semi-quantitative risk analysis employs a scoring scheme, such as a simple numeric scale ranging from 1 to 3, 1 to 5, or 1 to 10.

14. C. Residual risk is the risk that remains after risk treatment (whether mitigation, transfer, or avoidance) has been completed.

15. A. The best way to manage residual risk is to include it in the risk register and then subject it to the usual risk analysis, risk scenario development, and risk treatment.

16. C. Risk identification focuses on finding threats, vulnerabilities, and assets that could lead to loss events. Analysis and evaluation follow to determine the likelihood, impact, and overall risk level.

17. B. Semi-quantitative analysis bridges the gap between qualitative and quantitative methods by using numeric scales to consistently prioritize risks, without having to rely on detailed financial or statistical models.

18. B. A BIA identifies essential processes, resources, interdependencies, and the potential business impacts of downtime, helping prioritize systems and processes for protection and recovery planning.

19. B. The threat landscape reflects the external and internal threat environment; changes in attacker capabilities, techniques, or targeting patterns directly alter the organization's threat landscape.

20. B. Threat modeling systematically examines potential threat actors, attack vectors, and vulnerabilities to anticipate how an attack could occur and the potential damage that could result. It is a core risk identification technique preceding vulnerability testing and risk analysis.



Chapter 3

Risk Response and Reporting

This chapter covers CRISC Domain 3, "Risk Response and Reporting," including the Risk Response, Control Design and Implementation, and Risk Monitoring and Reporting subdomains. The job practice for this domain is described as follows:

This domain deals with the development and management of risk treatment plans among key stakeholders, the evaluation of existing controls and improving effectiveness for IT risk mitigation, and the assessment of relevant risk and control information to applicable stakeholders.

The job practice consists of these elements:

✔ A—RISK RESPONSE

1. *Risk Treatment/Risk Response Options*
2. *Risk and Control Ownership*
3. *Third-party Risk Management*
4. *Issue, Finding, and Exception Management*
5. *Management of Emerging Risk*

✔ B—CONTROL DESIGN AND IMPLEMENTATION

1. *Control Types, Standards, and Frameworks*
2. *Control Design, Selection, and Analysis*
3. *Control Implementation*
4. *Control Testing and Effectiveness Evaluation*

✔ C—RISK MONITORING AND REPORTING

1. *Risk Treatment Plans*
2. *Data Collection, Aggregation, Analysis, and Validation*

3. *Risk and Control Monitoring Techniques*
4. *Risk and Control Reporting Techniques (heat map, scorecards, dashboards)*
5. *Key Performance Indicators*
6. *Key Risk Indicators (KRIs)*
7. *Key Control Indicators (KCIs)*

The CRISC Task Statements relevant to this domain focus on risk treatment and reporting. The CRISC Task Statements covered in this chapter are:

5. *Establish accountability by assigning and validating appropriate levels of risk and control ownership.*
11. *Review the results of risk analysis and control analysis to assess any gaps between current and desired states of the IT risk environment.*
12. *Facilitate the selection of recommended risk responses by key stakeholders.*
13. *Collaborate with risk owners on the development of risk treatment plans.*
14. *Collaborate with control owners on the selection, design, implementation, and maintenance of controls.*
15. *Validate that risk responses have been executed according to risk treatment plans.*
16. *Define and establish key risk indicators (KRIs).*
17. *Monitor and analyze key risk indicators (KRIs).*
18. *Collaborate with control owners on the identification of key performance indicators (KPIs) and key control indicators (KCIs).*
19. *Monitor and analyze key performance indicators (KPIs) and key control indicators (KCIs).*
20. *Review the results of control assessments to determine the effectiveness and maturity of the control environment.*
21. *Report relevant risk and control information to applicable stakeholders to facilitate risk-based decision-making.*

The topics in this chapter represent 32% of the CRISC examination.

The CRISC Task Statements relevant to this domain address several key activities in risk response and reporting, beginning with the collection of information about an organization's business and IT environments, and the identification of impacts to the organization's objectives and operations. This includes the identification and evaluation of threats and vulnerabilities, and establishing accountability by assigning risk and control ownership. An IT risk register documents the risk profile. Identifying risk appetite and

tolerance and conducting security awareness training help promote a risk-aware culture. IT risk assessments and risk analysis help identify control deficiencies, leading to risk responses, risk treatment plans, and the design and implementation of controls—all of which should be validated. Risk reporting includes the use and monitoring of key control indicators, key risk indicators, and key performance indicators. Assessments of controls and analysis of risks can help understand their maturity and effectiveness. New and emerging technologies should be evaluated to identify changes in risks. All of this should drive risk-driven decision-making.

In this chapter, I review the concepts that comprise CRISC Domain 3, which are focused on risk response and reporting. There is a natural progression from the risk framing, analysis, and assessment discussions in the previous chapters, which covered CRISC Domains 1 and 2, to the discussion of risk response in this chapter; you will see that many of the outputs from those processes lead directly to this chapter because an organization generally wants a meaningful response to any risk to itself and its systems and data. After discussing how risk responses are chosen, I discuss how to implement those responses, including designing, developing, and adjusting controls, and finally reporting on risk.

Risk Response

Every organization has risk. The leaders of smart organizations, however, have determined how they will appropriately address that risk in their business context. Will they accept all risks? Will they transfer some risks to a third party, such as an insurance company? These are some questions to consider in the risk response process, as discussed in this chapter.

I also discuss the major control frameworks within the field and how they are incorporated into risk response. You should note that risk is different for every organization; risk within an educational context is different from a financial context, which is different from a government context. Also, keep in mind that IT risks are different from business or mission-oriented risks. It's important to note that the organization's mission also shapes risk responses. This is why it's so important that organizations conduct a thorough risk analysis and understand and prioritize the response options that are right for them. You should also remember that the risks to the business, the business goals and objectives, and the systems supporting those functions should all be considered throughout the risk response process. It's easy as an IT professional to forget that the business objectives are the top priority, but to maximize efficiency and effectiveness (and to pass the exam), you should keep that in mind.

Risk Treatment/Risk Response Options

Having developed your understanding of risk, you should also understand the corresponding risk response options. These options will help you bring your risk down to a level of risk tolerance your organization can accept. Understand that this is not just merely responding

to risk; *risk treatment* and response options should also meet the needs of the mission, business goals and objectives, and overall organizational strategy. A risk response option focused solely on reducing risk in a vacuum may interfere with normal business operations or reduce the effectiveness of the business when performing its mission. Therefore, you should consider risk responses not just in terms of reducing risk; they must also reduce risk without adversely impacting business processes. Within this context, consider how you would recommend that your organization respond to risk effectively, with a balanced approach. The next few sections discuss the different risk response options.

Risk Mitigation

Risk mitigation generally involves the application of controls that lower the overall level of risk by reducing the vulnerability, the likelihood of a threat exploiting that vulnerability, or the impact to the asset if the risk were to be realized. Mitigation actions take many forms; a policy requiring nightly offsite backups of critical data can mitigate the risk of data loss, for instance. Similarly, applying a newly released web browser patch can mitigate the risk of exploitation. The goal is to reduce the risk to a level acceptable to the organization's leadership. I discuss controls and their implementation in the "Control Design and Implementation" section, but for now, you should understand that you may have to add new controls, update, upgrade, or even change controls, or, in some cases, change the way your business processes work. All of these actions are designed to mitigate or reduce risk.

Risk Sharing

Risk sharing (often referred to as *risk transference*) entails the use of a third party to offset part of the risk. Risk sharing can involve outsourcing some activities to a third party to reduce the financial impact of a risk event. Sharing off-site (co-location) assets and contractual obligations with other entities is one way that organizations implement risk sharing; a cloud service provider can be used within this scenario. The cloud provider might be contractually obligated to assume part of the financial impact in the event of a breach, but be aware that there is a potential loss of brand goodwill or other intangible assets that can be difficult to offload.

Another example of risk sharing is the use of an insurance service to cover the financial impacts (at least partially) of a breach; however, the intangible losses mentioned previously would still be present. The point is that risk sharing (or transference) typically only covers a portion of the risk; it does not reduce all of the risk, nor does it transfer accountability or responsibility. Therefore, multiple risk response options will likely be needed concurrently.



Don't forget that the use of a risk-sharing scheme does not totally absolve an organization of its responsibilities; organizations may still retain responsibility—and, more importantly, accountability—if there is a loss of data or revenue. Additionally, legal liabilities may also be involved.

Risk Acceptance

There will always be some level of residual risk, no matter how many controls you implement or the amount of insurance you purchase. Again, the goal is to reduce the risk to a level acceptable to leadership. *Risk acceptance* occurs when the active decision is made to assume the risk (either inherent or residual) and take no further action to reduce it. The word *active* is critical here; risk acceptance differs from *risk ignorance* in that the risk is identified, alternatives are considered, and a conscious decision is made to accept it. A good example of this is when you have taken all practical steps to reduce the risk of an attack from a malicious entity. Your organization may have installed perimeter security devices, strong authentication methods, and intrusion detection systems, which will significantly reduce risk. However, since risk cannot be completely eliminated, there is always a chance that an attacker will enter the infrastructure through another means. This residual risk may have to be accepted because it might be financially unwise to invest more resources into further mitigating the risk. Similarly, a company might choose to willfully ignore government-mandated controls that would incur fines or other financial penalties, assuming the financial cost would be less than the cost of implementing the controls. As I discuss in "Evaluating Risk Response Options," the asset's value may be far less than the cost of implementing mitigating controls, so in those cases, risk acceptance may be the right choice.

Note that risk acceptance of any kind should be a formal management decision, fully documented, along with provisions for monitoring the risk, typically through the risk register. Some organizations will execute a *risk acceptance letter* to memorialize the decision and assign accountability. As with all risk response options, management should be cognizant of monitoring risk, since it frequently changes, and in turn make changes to the response as dictated by the changing operating environment, new technologies, and the threat landscape.



It's important to note that risk acceptance doesn't insinuate ignorance; it's an active acknowledgment and conscious decision to accept risk as it is, provided it falls within risk appetite and tolerance ranges.

Risk Avoidance

Risk can be avoided through several methods. Sometimes risk can be avoided by simply choosing not to participate in the activity that incurs the risk. For example, an organization may decide that a particular business venture or entry into a new market would incur far more risk than it is willing to take on. Perhaps you don't need to conduct the activity at all. Avoiding risk can mean eliminating the risk associated with a particular activity by not performing it. *Risk avoidance* is the option taken when, after all risk treatment options, including mitigating controls, have been considered, the level of risk remains unacceptable. An organization might choose not to take on a proposed project because of the probability of failure and subsequent loss of capital, for example. Another scenario might be the organization choosing not to adopt a cloud solution because of the sensitivity of the information that would be stored in the cloud. In any event, risk avoidance does not

mean simply ignoring the risk, as some people may believe. Risk avoidance requires careful thought and planning, as well as balancing the level of risk incurred in following a particular path versus eliminating the risk by not following that path at all.

Evaluating Risk Response Options

Risk mitigation, sharing, acceptance, and avoidance are all valid options you should consider using (sometimes simultaneously) to lower the risk you face. The key is to use these options together; you want to balance risk and reward to maximize the benefits. These benefits can be financial or other, perhaps difficult to quantify (such as morale, partnerships, or long-term strategy).

One of the key considerations in balancing risk options is the cost and effectiveness of any response option. For example, as briefly mentioned, if the value of the asset to the organization is much lower than the security controls required to protect it, is risk mitigation or reduction the right option? You should also consider the effectiveness of the controls; if all possible controls you could implement still don't reduce the risk to an acceptable level, are they worth it?

Other factors to consider in risk response options include those inherent to the organization itself. Sometimes organizational design and layout are a factor because of how different divisions and departments are structured within the organization and who they report to, which affects resource allocation, chain of command, and so on. Sometimes these factors must be changed to implement an effective risk response. Other considerations are governance, including legal and regulatory requirements, and organizational culture. An example of how organizational culture can affect a risk response is when an organization determines it needs to revoke administrative privileges from ordinary users.[1] While this response may not be expensive in terms of required equipment or additional controls, the cost is the shock to and pushback within the organizational culture. People normally do not like giving up what they already have, and administrative rights are one of those things that all users seem to believe that they need. Organizational culture will definitely impede the implementation of this particular risk response.

In evaluating and selecting risk responses, you should consider all the factors I've discussed previously, such as the cost of the response (and of the potential cost of not responding), organizational context and associated missions, risk tolerance, governance, and the simple ease of implementation.

Risk and Control Ownership

Risk and control ownership are issues that must be considered during a risk assessment and response. Depending on how the organization is structured and how the risk management strategy has been developed, risk ownership may be assigned to one or more managers

[1] End users having administrative privileges on their workstations is now considered a substandard, risky practice, as it increases the potential impact of malware, as well as increasing support costs when users change system configurations that cause system functionality and stability problems.

spanning multiple functional areas. This is because the risk that affects one area likely affects others as well, so many different people may be responsible for affected areas and required to deal with and respond to the risk.

Control ownership is also something that should be examined; often, controls that provide protection for a given asset or a number of systems may not fall under the operational purview of the risk owner. Some controls span the entire organization and protect multiple assets, rather than only a specific system. These are referred to as *common controls*, and the responsible entity for these controls is the common control's provider. For example, consider physical security controls. There may be a physical security officer for the organization who oversees guards, personnel security, and physical access to secure areas. They may also be responsible for locks, closed-circuit televisions (CCTVs), and physical intrusion detection systems and alarms. All of these controls protect information systems and data assets that may be in a secure data center. Risks may be presented to the person responsible for information systems in the data center, but the controls to protect them may fall under the physical security functional area. This would require coordination among the different functional areas if controls managed by one area are insufficient to protect assets managed by another. This may also affect organizational structures, budget, and allocation of resources.

It is important to identify and document risk and control owners in the risk register. This helps assign responsibility and accountability for risks, the responses, and the controls. Additionally, the risk and control owners should be given the authority to fulfill their responsibilities in reducing risk and to implement controls as needed. Both risk and control owners, as part of this authority, should have some level of control over budgets, personnel, equipment, and other resources under their purview. If owners are given the authority to carry out their responsibilities, they can be held accountable when risk is not properly managed or when controls are ineffective at reducing it.



Risk, asset, and control owners are not always the same person, functional area, or even organization. You must identify these owners early in the assessment process and maintain careful coordination and communication with them and other relevant stakeholders, within the boundaries of your authority and assessment scope. Having different types of owners can raise politically sensitive issues around resourcing, responsibility, accountability, and sometimes even blame.

Third-party Risk Management

In today's business world, an organization rarely functions independently from other entities. Most businesses are not entirely independent and self-sustaining; they depend on suppliers, service providers, and so on to help them bring their products and services to market. These third parties often perform services that are outsourced by the organization due to a lack of infrastructure or a desire not to perform these functions internally. A great example of a third-party service provider is one that provides cloud services to businesses, such as data

storage, applications, and even entire infrastructures, to organizations that don't have the means or desire to build the internal structures needed to manage them. Sometimes it is more cost-effective for an organization to simply hire a third party to perform a function or provide a service than to expend resources to create and maintain all the equipment and processes itself.

Risk Sharing Through Third Parties

If you recall the risk treatment options discussed earlier, transferring (or sharing) risk involves using third parties to bear some of the risk in performing certain business functions. Insurance is the classic example of risk transfer used in most risk management training texts, but this instance usually provides only monetary compensation in the event of a disaster or liability. Transferring risk to third parties also includes outsourcing certain functions and services, so this, too, is a standard risk treatment method. The organization, for example, can reduce its risk of accounting and budget errors with a small, untrained staff by outsourcing its accounting functions to a third-party accounting firm. *Third-party risk management* is the set of management activities used to identify and measure risks present in third-party organizations.

Accountability and Responsibility in Third-party Relationships

It's important to note that while this risk-management strategy is effective and common, the organization is transferring only risk, not responsibility. The organization almost always retains final accountability and responsibility for functions and services outsourced to a third party. For example, transferring a business's accounting functions to an outside firm mitigates risk by transferring some of it to the third party, but the organization is still ultimately responsible to its stakeholders, government regulators, banks, and financial obligations, even if the accounting firm makes an error. As another example, an organization that uses cloud data storage services still maintains accountability and responsibility for any data stored with a cloud provider if that provider suffers a breach that results in the release of sensitive information. While third-party providers may still incur some liability in these cases, the ultimate liability rests with the organization that contracted with them. Liability, accountability, and responsibility are also some of the key risk factors with third parties that I discuss next.

Third-party Risk Factors

As indicated earlier, several *risk factors* relate to the legal aspects of dealing with third-party providers. Most of these risk factors concern trust in the third-party provider or the organization's dependencies on the provider's services. For example, the organization must be able to trust that the third-party provider will provide adequate security for any data that is maintained in its infrastructure. In addition to trust, other risk factors include interoperability with existing systems in the organization, data protection and security, and legal liability. Remember, I said earlier that an organization can outsource risk to third-party providers, but the organization typically retains accountability, responsibility, and legal liability. Some legal liability, however, is imposed on the third-party provider, since it usually

has a duty to provide services and assurances that those services will meet certain legal standards, as defined in the contract between the provider and the organization.

The risk factors relating to third-party providers depend on the types of services provided and how the contract agreements are structured between the organization and the provider. Some contracts have cybersecurity provisions built into them. A *service-level agreement (SLA)* is the key to reducing risk when dealing with a third-party provider. The SLA stipulates what the third party must provide regarding service levels, performance, and so on. Some of these contracts are specific to the data type involved. For example, third-party agreements that govern protected health information, called *business associate agreements (BAAs)*, include cybersecurity provisions.

Another common approach is for companies performing business to respond to security questionnaires. There are standardized formats available on the Internet, but many companies rely on answers from security frameworks or design their own sets of questions to assess security/risk. More mature companies may conduct their own cybersecurity assessments, including audits, penetration testing, evidence gathering, and so on. In some cases, these can be just as intensive as an actual audit. All of these methods can serve as minimal cybersecurity standards for third-party providers. In many cases, they result in findings that need to be remediated to conduct business (meet minimal risk thresholds).

External Ratings and Continuous Third-party Monitoring

Because many organizations tend to outsource a majority of their IT stack, they are increasingly relying on external, continuously updated cybersecurity rating services to gain insight into the security posture of their vendors, service providers, and other third-party partners. These services, offered by firms such as SecurityScorecard, BitSight, RiskRecon, UpGuard, and others, evaluate the observable attributes of a vendor's external attack surface and convert those observations into a standardized security score. Although each rating provider uses its own scoring methodology, the general approaches are similar: data is collected from Internet-facing systems, threat intelligence feeds, malware sinkholes, DNS records, leaked credential databases, and other sources that reflect how well the organization manages certain aspects of its cyber hygiene.

These rating services are not a replacement for contractual due diligence, audits, or security questionnaires, but they do provide an independent, consistently measured signal that can be used to track changes in certain aspects of risk over time. For example, a sudden decline in a vendor's rating may signal emerging operational issues (such as patching delays, exposed services, unresolved vulnerabilities, or increased malicious activity targeting the vendor's infrastructure), which could be caused by financial distress or staffing issues. When used correctly, rating trends can complement traditional oversight methods by highlighting where deeper inquiry or follow-up is needed.

It is important to recognize the benefits and limitations of these ratings. On the positive side, ratings offer continuous visibility, require no direct integration with the vendor's internal systems, and allow comparisons across large supplier ecosystems. They help organizations detect risk conditions that might not be disclosed during annual assessments or periodic audits. However, ratings reflect only what is externally observable; they do not

capture internal control strength, detailed configurations, or the operational maturity of internal security functions. Additionally, these ratings, much like credit scores, are only a point-in-time rating and don't always reflect the longer-term cybersecurity posture of the target organization. A high rating does not guarantee that internal processes are sound, and a low rating does not always indicate systemic risk; it may simply reflect temporary misconfigurations or incorrect asset attribution.

Risk practitioners typically use external ratings in three ways:

- **Screening new suppliers** before onboarding, especially when handling sensitive data
- **Monitoring existing vendors** for changes that could indicate emerging risk or deterioration in control performance
- **Prioritizing assessments** to determine which suppliers require additional attention, attestation, or contractual follow-up

When these services are integrated into third-party risk management workflows, they can improve visibility without imposing additional burden on suppliers. However, any actions based on rating changes should be validated through direct communication, as ratings alone do not provide the context needed to determine the severity or business relevance of a particular observation.



Third-party risk rating services should augment, not replace, other aspects of third-party due diligence, such as contract review, questionnaires, and external attestations.

Third-party Threats

As with the risk factors mentioned earlier, threats associated with third-party providers primarily relate to trust and the contractual and legal aspects of those arrangements. Threats include failure to provide services as specified in contracts, failure to meet performance or security levels, failure to protect data, and loss of availability for systems or data to an organization's users. Third-party providers can certainly face some of the same threats as the organization, such as natural disasters. Therefore, some of the threats could also include loss of service or data for the business in some of these negative events. Third-party providers are also susceptible to hacking and malicious insiders, just like any organization. Therefore, threats of data loss or disclosure are also present with third-party providers.

Third-party Vulnerabilities

Vulnerabilities that result from associations or dealings with third-party providers include failure to adequately address contingencies in contractual agreements, such as system or data availability or loss, lack of or inadequate security protections, and service levels that fall below an acceptable standard. Furthermore, any vulnerabilities inherent in the third-party provider's organization are, unfortunately, indirectly incurred by any organization that deals

with the third party, since those vulnerabilities could affect data in its custody or services provided to others. Examples of internal vulnerabilities that could affect organizations for whom the third party provides services include infrastructure vulnerabilities (such as patch or configuration management issues, loss of equipment, or malicious intrusion by hackers), organizational vulnerabilities (such as loss of revenue or bankruptcy), and even technological vulnerabilities (such as a lack of sufficient encryption strengths or strong authentication methods). All these vulnerabilities could indirectly affect any organization that has a contractual agreement with a third-party provider.



Remember that even if risk is transferred to a third party, accountability and responsibility usually remain with the business.

Cloud Shared Responsibility Models

Organizations increasingly rely on cloud service providers (CSPs) to support critical business operations, making cloud security a central component of both risk treatment and third-party risk management. Cloud environments use a *shared responsibility model*, which defines the division of security and operational duties between the CSP and the customer. Understanding these boundaries is essential for selecting appropriate controls, validating risk treatment options, and determining what accountability remains with the organization.

The shared responsibility model varies slightly among cloud providers, but the underlying principles are consistent across service models: *Infrastructure-as-a-Service (IaaS)*, *Platform-as-a-Service (PaaS)*, and *Software-as-a-Service (SaaS)*. A typical cloud responsibility model is depicted in Figure 3.1.

FIGURE 3.1 A typical shared responsibility model.

<table>
<tr><th>Party</th><th colspan="4">Responsibilities</th></tr>
<tr><td rowspan="4">CUSTOMER</td><td colspan="4">CUSTOMER DATA</td></tr>
<tr><td colspan="4">Applications, Tools, Users, Certificates</td></tr>
<tr><td colspan="4">INFRASTRUCTURE SERVICES</td></tr>
<tr><td colspan="2">Data Encryption and Protection</td><td colspan="2">Network Encryption and Protection</td></tr>
<tr><td colspan="5"></td></tr>
<tr><td rowspan="4">CLOUD PROVIDER</td><td colspan="4">SOFTWARE</td></tr>
<tr><td>Operating Systems</td><td>Data Storage</td><td>Database Management Systems</td><td>Internet Connectivity</td></tr>
<tr><td colspan="4">HARDWARE</td></tr>
<tr><td colspan="4">Regions, Zones, and Data Centers</td></tr>
</table>

Concept of Shared Responsibility

In a shared responsibility model, the CSP is responsible for securing the cloud infrastructure, while the customer is responsible for securing data, identities, and configurations within the cloud. Responsibilities shift depending on the service model, but *accountability for organizational risk always remains with the customer*, regardless of how responsibilities are divided.

This distinction is important, because misinterpretation of shared responsibility often leads to unmanaged or misunderstood risk, especially in areas such as identity management, data protection, configuration, and monitoring.

Division of Responsibilities by Service Model

The allocation of responsibilities varies somewhat between the three principal cloud service types (IaaS, PaaS, SaaS), as explained next.

Infrastructure-as-a-Service (IaaS)
This model provides the customer with the most control, but also the greatest responsibility.

- The CSP manages:
  - Physical data center components
  - Hypervisors
  - Underlying network layers
- Customers manage:
  - Operating system configuration and security
  - Network architecture, segmentation, and firewall rules
  - Identity and access management for administrators and end users
  - Encryption of data at rest and in transit
  - Application security
  - Logging, monitoring, and incident response activities within the hosted environment

Platform-as-a-Service (PaaS)
The CSP manages the infrastructure layer (physical data center components, hypervisors, and network) plus runtime environments, middleware, and managed databases.

- The CSP manages:
  - Physical data center components
  - Hypervisors
  - Underlying network layers
  - Software development, maintenance, and security
- Customers manage:
  - Application configuration
  - End-user access management, access controls, and authentication mechanisms

- Data classification and protection
- Monitoring and user behavior analytics

PaaS reduces operational overhead, but requires the organization to manage data security controls and configuration decisions.

Software-as-a-Service (SaaS)
The CSP manages nearly the entire technology stack, including applications, platforms, infrastructure, and supporting services.

- The CSP manages:
  - Physical data center components
  - Operating systems and hypervisors
  - Underlying network layers
  - Software development, maintenance, and security
- Customers manage:
  - End-user access and identity governance
  - Data classification and retention
  - End-user security settings
  - Integration with downstream systems
  - Legal and regulatory compliance obligations

Even in SaaS environments, the customer remains responsible for preventing unauthorized access, ensuring appropriate use, and governing sensitive data.

Implications for Risk Ownership and Control Selection

The shared responsibility model directly affects both risk ownership and control design. Specific implications include:

- **Identity and access management** remains the customer's responsibility across all service models, making it a central control for cloud risk mitigation.
- **Misconfigurations** remain one of the greatest causes of data exposure and breaches, reflecting the customer's responsibility for securing logical configurations.
- **Data protection**, including encryption, retention, and classification, always requires customer decisions, regardless of the CSP's default capabilities.
- **Monitoring and logging responsibilities** vary by service model, but customers must ensure visibility into events relevant to their workloads and regulatory obligations.
- **Incident response** is shared; the CSP handles infrastructure events, while customers must manage application-level and data-level incidents.
- **Compliance alignment** remains the customer's obligation, including verifying that the CSP meets required certifications and that internal (customer) processes address remaining gaps.

These responsibilities must be reflected in risk treatment plans, control selection decisions, and ongoing monitoring.

Governance and Contractual Considerations

Cloud risk cannot be managed solely through technical controls. Governance requirements include:

- Formal documentation of shared responsibility matrices for each service model in use
- Review of CSP certifications (such as SOC 1, SOC 2, ISO/IEC 27001, CMMC, and FedRAMP)
- Verification of data residency, data retention, and deletion processes
- Definition of logging and monitoring responsibilities at all layers
- Contractual requirements for breach notification, audit rights, service levels, and termination
- Assessment of multitenancy risk and any compensating controls and attestations provided by the CSP

These governance elements must be included in third-party risk assessments and reflected in risk registers.

Continuous Oversight

Because customer requirements, cloud features, regulations, and CSP services change frequently, shared responsibility must be continuously reviewed. The organization should periodically reassess responsibilities, validate inherited controls, and ensure new features or changes do not introduce unexpected risk.



Remember that organizations can outsource many aspects of their operations, but they cannot outsource accountability.

A financial services firm deployed a new customer analytics platform built on a SaaS solution provided by a large cloud vendor. The implementation relied heavily on built-in vendor security features, but the firm did not assign internal control owners for identity governance, access reviews, or data retention settings. When the cloud provider modified its default logging-retention policies during a platform upgrade, the firm lost access to audit logs needed for a regulatory inquiry.

The issue was traced to a misunderstanding of the shared responsibility model: the cloud provider managed infrastructure and application-level controls, but the customer was responsible for retaining logs, managing identities, and configuring data governance policies. The event prompted the firm to update its vendor onboarding checklist, require explicit shared responsibility matrices, and implement continuous third-party monitoring using external rating services to track changes in vendor posture over time.

Third-party Risk Management

Managing third-party risk can be challenging and is an ongoing process. The initial risks may change due to a growing relationship with a business partner, additional data being sent to a third party, or changes in the threat landscape. Considerations for managing that risk may be performed annually, as contracts and technology change, and so on. As mentioned previously, some assessments are based solely on questionnaires, but other tools may provide ongoing technical assessments of third parties. These can help detect whether the effectiveness of a third party's security controls is trending in a negative direction, which, in turn, can be used as additional information to more fully inform risk assessments of third parties.

Business Process Outsourcing

Third-party risk management must account for *business process outsourcing (BPO)*, not just technology vendors. BPO arrangements involve external organizations performing operational or administrative functions that are integral to business outcomes. Common examples include payroll processing, call centers, claims administration, customer support, IT desktop support, accounting, human resources, and procurement services.

Although BPO providers may rely on technology, the primary risk exposure often arises from people, processes, and operational/procedural dependencies rather than software or infrastructure. Risks can include errors in transaction processing, failure to meet service-level expectations, regulatory non-compliance, data-handling weaknesses, labor disruptions, and reputational harm from poor customer interactions. In regulated industries, outsourced processes may also create compliance and accountability challenges if roles, responsibilities, and oversight mechanisms are unclear.

Effective third-party risk management requires identifying outsourced business processes, assessing their criticality to operations, and evaluating the provider's internal controls, governance practices, and resilience capabilities. Contracts should clearly define performance metrics, audit rights, data protection obligations, and escalation procedures. Treating BPO providers as part of the extended enterprise helps ensure that operational and non-technology risks are appropriately identified, monitored, and managed.

Issues, Findings, and Exception Management

Rarely do our security processes and risk response efforts go as smoothly as we would like. There are always things that don't go as planned. Sometimes there are conflicts between security requirements, competition for resources within the organization, findings that cannot be easily mitigated, and other issues and exceptions that must be addressed. These departures from planned activities and from established configuration baselines are cause for concern, especially when we must implement solutions that fall outside the normal security parameters we would like. Security and risk, like most things in life, involve compromise and balance. As security and risk practitioners, we are always trying to balance security, functionality, and resources, which seem to be in constant conflict. Security and functionality are almost always inversely proportional; for example, the

more secure a system is, the less functional it is for its users, and vice versa. Resources are always at a premium, and we never have enough money, qualified personnel, equipment, or facilities to achieve the level of security we would like. Additionally, risk itself is a give-and-take. We are often forced to accept a level of risk based on how cost-effective it is to implement controls to reduce the risk to a manageable level. In any case, how we deal with departures from the normal baselines, in the form of issues, findings, exceptions, and so on, is critically important.

Change and Configuration Management

One of the primary sources of issues with our security posture, in the form of findings and exceptions, is our change and configuration management processes and status. Remember that *change management* is a formal management process designed to ensure that needed changes are carefully requested, considered, approved, tested, and implemented in a safe, secure, and predictable manner. When implemented properly, the change management process should allow an organization to make major infrastructure changes without unexpected or critical consequences. However, change management can cause issues when emergency changes must be made for security or other reasons, and there simply isn't enough time to route the change through the organization's change control board (CCB). There are also issues if people do not follow the steps of a proper change management process and simply implement changes on the fly without approval. Both issues should be carefully considered because they could incur additional risk, represent an undesirable departure from the normal change process, and must be minimized whenever possible.

Configuration management brings the same challenges. While change management applies to major infrastructure changes, configuration management is a subset of that process and is specific to individual systems. Arbitrarily changing configurations without careful planning and testing could result in unstable or unusable systems. If emergency configuration changes must be made, this represents additional risk because there may not always be time for the formal testing and approval process. Configuration management is also affected by legacy systems that cannot be updated if a critical security issue is discovered, such as a zero-day vulnerability. These problems could crop up during a vulnerability scan, where a vulnerability is discovered, and the patch either does not exist or cannot be applied to an older system, or to a system that, if patched, would break line-of-business systems. This is where an exception would have to be considered: the company could either allow the system to operate under its current configuration, while acknowledging the security risk, or risk the impact of the update on the line-of-business systems, resulting in a loss of potential revenue. Even when the risk is considered and an exception must be made, this, again, is an undesirable departure from the normal process steps. I discuss exception management in more detail in the next section.

Release management is also a subset of change management and, along the same lines, entails risks when software is released ahead of schedule or before it is adequately tested and assessed for risk. The release management process must, as much as possible, adhere to scheduled releases that are tested and include risk-reduction measures. As discussed in the preceding paragraphs, emergency releases entail risk and should be avoided whenever possible.

Exception Management

No matter how standardized security controls are, there will always be exceptions, and you need a process to manage them and keep them to a minimum. Consider the example in the previous section where an operating system patch cannot be applied because of a legacy application that is critical to the organization's mission. This mitigation is important, but less important than the legacy application and its associated mission. *Exception management*, as mentioned earlier, comes into play here; you should evaluate the risk of implementing the patch and breaking the application, thus putting the organization's mission in jeopardy, versus the risk of not implementing the patch and dealing with the security ramifications that come with that decision. Management will have to decide which is the riskier choice. If the choice is to make an exception and not implement the patch, you need to document this as an accepted exception to the policy and track it, usually in the risk register. Perhaps the patch can be implemented in the future, or another mitigation could arise that makes the exception moot. This is another aspect of risk acceptance: you are asking management to accept an ongoing risk as an exception because the response cannot be implemented for whatever reason. Either way, keeping a close eye on the issue is key to not letting these exceptions slip through the cracks.

Regardless of where issues, security findings, or exceptions to normal processes, policies, or baselines occur, these departures should be avoided as much as possible; when unavoidable, they must be carefully considered, with great care taken to reduce the risks they present.

Workflow for Issues, Findings, and Exception Management

An effective risk management program requires a structured, repeatable workflow for identifying, evaluating, approving, tracking, and resolving issues, findings, and exceptions. These events represent deviations from required security baselines, policies, or control expectations. Without a standardized workflow, organizations risk allowing unresolved weaknesses, unmanaged residual risk, and inconsistent decision-making across business units. The following suggested workflow outlines the recommended steps for managing issues, findings, and exceptions. This workflow should be tailored as needed to harmonize with existing policies and practices.

Step 1. Identification and Logging

The workflow begins when an issue, finding, or exception is identified. Sources may include audits, assessments, vulnerability scans, change management activities, incident response, operational monitoring, or threat intelligence. Each identified item must be logged in a centralized repository or governance, risk, and compliance (GRC) system. The log entry should include:

- A clear description of the issue or deviation
- Date of discovery
- Source of identification
- Affected systems, processes, or controls
- Initial severity or classification
- Related artifacts that substantiate the item

Centralized risk issue logging ensures visibility and traceability throughout the lifecycle of an application or system.

Step 2. Initial Triage and Classification

Once logged, each item undergoes an initial triage to determine urgency and potential risk. Classification criteria typically include:

- Likelihood of exploitation or recurrence
- Potential business impact
- Regulatory or contractual obligations, such as incident notification
- Whether the deviation is temporary or systemic

Triage assigns severity levels (e.g., critical, high, moderate, low), which guide prioritization, communication, escalation, and required review timelines.

Step 3. Risk Assessment and Impact Analysis

A detailed risk assessment is performed to evaluate the consequences of leaving the issue unresolved. This step examines:

- Associated threats
- Exploitable vulnerabilities
- Business impacts (financial, operational, reputational, legal)
- Compliance implications
- Potential cascading or downstream effects

The analysis may draw on existing risk scenarios or create a new one if the deviation represents an unmodeled risk. The retention of historic risk assessments aids in this research.

Step 4. Determination of Required Action

Based on the risk assessment from the previous step, the organization categorizes the item:

- **Issue:** A deviation or problem that requires corrective action.
- **Finding:** A formally reported deficiency from an audit, assessment, or compliance activity.
- **Exception:** A temporary, formally approved deviation from a required policy or control.

For issues and findings, *corrective actions* or *preventive actions* must be proposed. For exceptions, a justification must be documented that explains why the control cannot be implemented as required.

Step 5. Review and Approval

Items must be reviewed by the appropriate authority, typically:

- Cybersecurity or risk management leader (to facilitate these steps)
- The control owner (technical evaluation)
- The risk owner (business impact evaluation)
- Compliance, internal audit, or security leadership (policy/regulatory alignment)

Exceptions require formal approval by management. Approval must specify:

- Duration (time-bound expiration)
- Required compensating controls
- Monitoring and reporting requirements
- Conditions for reassessment or revocation

No exception should be open-ended or permanent unless explicitly approved through a governance process.

Step 6. Action Plan Development

For issues and findings, a *risk response action plan* is developed. This includes:

- Specific remediation steps
- Assigned owners
- Required resources
- Target completion dates
- Success and acceptance criteria

For exceptions, the plan documents compensating controls and the activities required to monitor the residual risk. The action plan should be reviewed and approved by the appropriate personnel before implementation.

Step 7. Implementation and Tracking

All actions must be tracked to completion. Tracking includes:

- Monitoring progress against deadlines
- Documenting changes or delays
- Updating affected risk scenarios or controls
- Ensuring resources remain available for remediation

A GRC tool or centralized register is typically used to maintain visibility into status.

Step 8. Validation and Verification

Once remediation is complete or an exception is nearing expiration, the organization validates that the item has been resolved appropriately. Validation activities may involve:

- Technical testing
- Control observation
- Documentation review
- Reassessment of the risk scenario
- Audit confirmation

For exceptions, verification ensures compensating controls remain effective and that a reevaluation occurs before expiration.

Step 9. Closure and Documentation

The item is formally closed when the following are completed:

- All actions are confirmed to be completed
- Validation results confirm effectiveness
- Residual risk is documented and accepted
- Leadership approves closure (where required)

All closure information is retained for audit, governance, and reporting.

Step 10. Reporting and Continuous Monitoring

Issues, findings, and exceptions contribute to the organization's overall risk posture. Summary information must be reported periodically to stakeholders, including:

- Counts by context and severity level
- Aging metrics
- Trends by business unit or control area
- Repeated or systemic deficiencies
- Exceptions nearing expiration

Continuous monitoring ensures that emerging patterns are detected, unresolved items do not accumulate, and the organization maintains compliance with governance requirements.

Step 11. After-action Review

Organizations may choose to conduct an *after-action review (AAR)* on significant matters, to ensure that processes, tools, personnel, and outcomes for future issues are managed and resolved effectively. Root-cause analysis may be warranted to better understand difficulties encountered.

Management of Emerging Risk

Risk is not static; it is constantly changing. The dynamic nature of risk means that it must be constantly monitored, reassessed, and reevaluated, and risk responses will often need to change as risk changes. The components of risk, including threats and vulnerabilities, also change. The organization itself changes as it is restructured, as its mission and goals evolve, as its place in the marketplace evolves, and, of course, as its risk appetite and tolerance change. Technologies also change, and organizations sometimes are in a hurry to adopt the latest and greatest technology because they are convinced they need to; sometimes that need is indeed valid, but sometimes the technology marketing machine convinces an organization it has a need where before it did not. Emerging technologies and their associated threats, vulnerabilities, risks, and, by extension, risk responses must be continually evaluated.

Emerging Technologies

New or emerging technologies can pose risks to organizations, as there are several important considerations when integrating the latest technologies into existing infrastructure. Many organizations have an unfortunate tendency to rush out, buy, and attempt to implement

the latest (and often unproven) technologies on the market, without careful planning or consideration of the existing infrastructure and how it will react to the new technologies.

One of the primary considerations organizations must address is making a business case for a new technology. This case may be to fill a gap that the older technology does not provide or to provide a capability the organization must have to compete in the market. In other words, the company must really justify new technologies to make them worth the risks they could incur.

At the opposite end of the spectrum, many companies are averse to investing significant resources in new technologies, instead relying on older, tried-and-true technologies that have always worked. These organizations fail to consider the changing threat landscape, and sometimes the need for maintaining a cutting edge in the marketplace. The threat landscape can render older technologies obsolete and ineffective in protecting the company's information assets. The market a particular organization participates in can also drive the need to update its older technologies to those that enable it to produce goods and services faster, more efficiently, and to better meet the needs of its customers.

In any event, emerging technologies present challenges to organizations that must be considered. These challenges include the proliferation of AI, emerging risk factors, vulnerabilities, and threats associated with those technologies. I discuss those in the next few sections.

Emerging Risk Factors

Emerging technologies pose several risks inherent to their integration into existing infrastructure. If the organization has justified the implementation based on a true business need for the emerging technology, it must consider several risk factors. One of the major risk factors is interoperability with existing infrastructure. Frequently, newer technologies don't work properly with older systems right out of the box; adjustments to the existing infrastructure may be needed to integrate the new technology, or bridging technologies may be required to connect the two. Interoperability doesn't just involve the right connections; it also involves data formats and flows, security methods, and interfaces to other systems. These considerations, and many others, are risk factors that must be considered before acquiring and integrating new technologies.

Another risk factor is security. New technologies may have security mechanisms that are not necessarily backward compatible with existing ones. Examples include encryption algorithms and strengths, identification and authentication technologies, integrity mechanisms, and even redundant or backup systems. Additionally, the systems may involve a learning curve that might intimidate users who must now learn new security methods for the system. The human factor can be a weak link in the security chain, so either a lack of training or a lack of adaptability to the new security mechanisms can introduce risk.

System updates and changes can also be risky if not managed properly. Integrating new technologies into an environment with older ones can introduce both intentional and unintended changes, affecting the stability of the organization's implementation, operations, and maintenance of a particular system, making change a risk factor. Even in the disposal or replacement phase of a system's lifecycle, introducing new technologies to replace older ones can be problematic if not planned and executed properly. New technologies that are not

adequately tested for functionality, performance, integration, interoperability, and security may not be able to replace older systems effectively, resulting in extended costs and possibly requiring the extension of the older systems' lifecycle.

Emerging Threats

Threats resulting from the introduction of new and emerging technologies into the existing infrastructure are numerous. If the organization has failed to adequately plan for the new technology, these threats can become significant. These include untested or unproven technologies, non-interoperability, incompatible security mechanisms, and a lack of suitability of the technology for use in the organization. The organization could also incur additional costs and require more resources because of a faulty implementation. These threats can be minimized by careful planning and integrating new technologies using a stable systems development lifecycle (SDLC) model.

Additionally, organizations should implement both threat modeling and threat hunting processes into their security program when new technologies are considered. Threat modeling and threat hunting can help an organization identify not only the generic threats that affect all organizations, but those specific to its own business model, processes, assets, vulnerabilities, and risk scenarios involving the new technology.

Emerging Vulnerabilities

As with threats, the vulnerabilities associated with working with emerging technologies are numerous. General vulnerabilities could include a lack of trained staff committed to managing and implementing the new technology. Inadequate project planning is also a serious vulnerability that could affect the organization's ability to effectively integrate new technologies. Another vulnerability could be a weak support contract or other type of warranty, guarantee, or support for a new technology or system. Most of these vulnerabilities appear when new technologies are first implemented and tend to be mitigated or reduced as the technologies are integrated into the existing infrastructure, but they still persist.

Beyond general vulnerabilities, new technologies also have their own inherent technical vulnerabilities. Often, these vulnerabilities are not discovered until long after a technology has been implemented. Although vendors are increasingly security-conscious in implementing and designing technologies from a secure perspective, zero-day vulnerabilities are often discovered months or even years after those technologies are installed and operational.

To make matters more challenging, some vulnerabilities are not detectable by vulnerability management tools. They are only detailed by the vendor, and only if they happen to detail them. In some situations, this makes monitoring application updates even more relevant, including in the vulnerability management process.

When implementing a new technology, organizations should make an extra effort to conduct vulnerability scanning and assessments, and pay particular attention to new vulnerabilities, even those that seem unrelated to the new technology. Often, these vulnerabilities are indirectly related and arise from interoperability or integration issues with the infrastructure.

In some cases, it may make sense to perform additional validations against new technology—especially if there are additional risk factors, such as large amounts of sensitive data. Validations may include, but are not limited to, penetration testing, secure configuration reviews, and application security testing.



The key areas of concern with emerging technologies are interoperability and compatibility. These two concerns affect the security, functionality, and performance of new technologies when installed in existing infrastructure.

AI and Machine Learning as Emerging Risk Domains

Artificial intelligence and machine learning (AI/ML) systems represent a rapidly evolving class of technologies whose risk characteristics differ in important ways from traditional software or infrastructure. Unlike conventional systems, where functionality is defined by static rules, predictable upgrade cycles, and well-understood dependencies, AI/ML models' behavior changes, based on training data, tuning processes, and model evolution. These characteristics introduce a new category of emerging risk that organizations must evaluate carefully as adoption expands.

One of the more fundamental challenges is that AI/ML systems can behave unpredictably when exposed to data or conditions they were not trained to handle. A model that performs reliably during testing can fail in production when presented with inputs that differ only slightly from its training distribution. In cybersecurity contexts, adversaries can intentionally craft such inputs to trigger unexpected behavior, including a class of vulnerabilities known as adversarial attacks. These weaknesses are often invisible during traditional control assessments because the model appears to function normally under expected conditions.

AI/ML systems also inherit risks from the data used to train them. *Bias*, omissions, or inaccuracies in training data can directly influence outputs and decisions, leading to operational errors, hallucinations, compliance violations, or fairness concerns. From a risk-response perspective, the problem is not simply that the model is "wrong," but that the error is embedded deep within data pipelines or tuning processes that may not be transparent to technical reviewers.

Model drift is another emerging risk. Over time, AI/ML models lose accuracy as business conditions, user behavior, threat environments, or operational contexts shift. Without monitoring, this drift can continue unnoticed until the model's output no longer aligns with business requirements or the organization's risk appetite. Monitoring drift requires specialized metrics and continuous assessment, which differ substantially from traditional system monitoring workflows.

AI/ML systems also expand the *attack surface*. Models can inadvertently expose training data through membership inference attacks or reconstruction techniques, creating privacy risks, even when conventional data protection controls appear sufficient. Dependencies on third-party model repositories, pre-trained embeddings, or cloud-hosted model APIs introduce additional supply chain risks, especially when those components are not vetted or their development processes are opaque.

Another often overlooked issue is *explainability*. Many modern AI/ML systems produce outputs that are difficult to interpret, making it challenging for risk practitioners to determine whether a model's behavior aligns with governance requirements. This lack of transparency complicates control testing, reduces auditability, and may limit the organization's ability to justify risk-based decisions to regulators, stakeholders, or customers.

Finally, the pace of AI/ML innovation itself is a source of risk. New techniques and model types emerge rapidly, and many organizations adopt them before suitable controls or assurance practices are fully developed. This "capability-first, security- and governance-later" pattern is a recurring contributor to emerging risks across industries.

For these reasons, AI/ML systems require differentiated risk identification and treatment strategies. Traditional technical controls remain important, but they must be supplemented with model governance practices, including data provenance tracking, drift monitoring, robustness testing, access controls for training pipelines, version management, and periodic evaluation of model accuracy and fairness. As AI/ML adoption continues to expand, organizations must incorporate these considerations into their broader emerging risk processes to ensure that model behavior remains aligned with organizational objectives and risk tolerance.



For a complete examination of AI-related risk and governance, refer to the *AIGP Study Guide* (Wiley Publishing) by this book's author.

Control Design and Implementation

A *control* is a measure or safeguard put into place to protect organizational assets (including systems, equipment, data, facilities, processes, and even people), ensure compliance with laws and regulations or other governance requirements, and reduce overall risk to the organization and its assets. Controls can help improve an organization's overall security posture by addressing a gap or making up for a shortfall in another control. Controls can be technical solutions, such as firewalls or intrusion detection systems, or processes, procedures, policies, or physical controls.

Control Types, Standards, and Frameworks

In this section, I discuss the various types and functions of controls as well as different control frameworks that can help an organization secure its assets. In later sections, I also discuss how an organization can design, select, and analyze its controls for effectiveness. Additionally, I cover the specifics of implementing controls and control testing. So, how do you determine which controls to apply to your system, or whether your existing controls are sufficient? I discuss this in more depth in the following sections.

Types of Controls

As you may have gathered by now, there are many types of controls to choose from. However, you can generally group these controls into broad categories by their intended type and function. Depending on which control framework you read, or which cybersecurity or risk expert you are talking to, these controls may be referred to differently, as some standards categorize them in different ways. However, for our purposes and for the CRISC exam, you need to understand three main types of controls: administrative (also called managerial), technical (also called logical), and physical (also called operational). The following are examples of each:

Administrative (Managerial) Control

- Requiring an acceptable use policy to dictate IT usage within the organization
- Starting a risk management program to identify, mitigate, and monitor unacceptable risk levels within the organization
- Implementing additional policies and procedures within the organization to support external governance and the desires of management

Technical (Logical) Control

- Implementing additional firewalls to protect internal (non-public-facing) systems
- Adding 802.1X port-based network access control to deter unapproved assets from being added to the network
- Installing an intrusion detection system/intrusion prevention system (IDS/IPS) to monitor for malicious activities or violations of policy

Physical (Operational) Control

- Implementing physical intrusion detection and alarm systems throughout the facility
- Deploying a human-based guard force to patrol sensitive areas
- Construction of fencing, bollards, and other physical obstacles to prevent unauthorized entry into an organization's campus

Controls are also categorized by function—that is, what purpose they serve and how they work. The most common control functions are deterrent, preventive, detective, corrective, compensating, and recovery. Table 3.1 lists each control type, its function, definition, and an example of the control.



Controls may be categorized by function: deterrent, preventive, detective, corrective, compensating, and recovery. Often, controls can serve more than one function simultaneously, making it difficult to categorize them into a single category.

Note that different control functions can fall into one or more control types: deterrent, preventive, detective, corrective, compensating, and recovery controls, all of which are technical, administrative, or physical. You will also notice some significant overlap between these types and functions. For example, video surveillance can be a deterrent control as well as a detective control.

TABLE 3.1 Control functions and descriptions.

Control Function	Definition	Example
Deterrent control	Deters individuals from performing policy violations or illegal acts	Video surveillance, guards, login warning banners, policies
Preventive control	Prevents policy violations or illegal acts	Firewall rules, permission settings, perimeter fencing
Detective control	Detects or discovers policy violations or illegal acts	Physical intrusion detection systems and alarms, system intrusion detection systems, integrity checks, audit logs, video surveillance
Corrective control	Restores normal operating state	Removal of malware and applying a clean system image to a system infected with malware
Compensating control	Longer-term alternative solution when primary controls cannot be implemented for legitimate business reasons; used to reduce risk in lieu of a desired control	Isolation and enhanced monitoring of a legacy system that does not support disk encryption
Recovery control	Restores availability and continuity of operations after a disruption or failure	Activation of a disaster recovery site and restoring data from backups after a data center outage

A deterrent control must be known to an individual or entity for it to be effective, since it dissuades them from committing a policy violation or an illegal act. For example, a video surveillance camera can be a deterrent control—if a person sees it in operation, they are less likely to commit a malicious act because they know they are being viewed and recorded. This deters them from that act. A preventative control, on the other hand, does not have to be known by an individual for it to be effective. An example of this would be a firewall rule that prevents a user from surfing to a malicious site. The user doesn't have to know that the firewall rule exists for it to still be effective and prevent them from committing this policy violation.

Another difference worth noting is the distinction between corrective and compensating controls. As indicated in Table 3.1, compensating controls are implemented to replace a preferred or primary control that can't be implemented for legitimate business reasons (e.g., lack of resources). Even if the primary control can't be implemented, the risk must still be reduced, so a control (or sometimes multiple controls) is implemented in its place. Compensating controls are longer-term controls that are unlikely to change in the immediate future.

Corrective controls, however, are temporary, short-term duration controls. They are designed to correct an immediate security issue. For example, let's say that a storm comes through your area and destroys a section of the strong concrete wall surrounding your facility. Now there is a huge hole in the perimeter, which could allow an unauthorized person to enter the facility grounds. To immediately mitigate this risk, you may post a guard to patrol the area for a few days until you can come up with a better solution. Obviously, keeping a guard there on a long-term basis is not wise financially or from a resource usage perspective. After a few days, you might be able to fix the wall, if possible, which restores the primary control to its operational state. If for some reason that portion of the wall can't be fixed, however, you may install temporary fencing or another physical obstruction in its place, and supplement it with surveillance cameras, an intrusion detection system, and additional alarms so that it would reduce the risk of an intruder entering the area, without having to keep a guard there all the time. These might be considered compensating controls, since the primary control (the strong concrete wall) can't be addressed immediately.

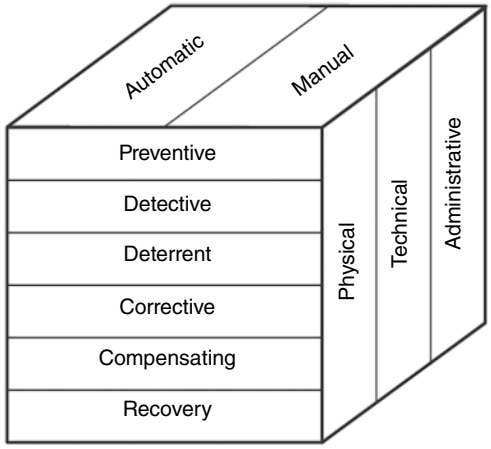
As mentioned previously, control types and functions are multidimensional and can overlap. They can also complement each other. For a given administrative control (e.g., a policy that dictates that all sensitive data must be encrypted), there is usually a technical control, such as the encryption mechanism, and possibly even a physical control, such as hardware encryption devices being locked in a secure area, to back it up. Again, the various standards and frameworks, which I discuss shortly, often categorize these types and functions of controls differently, so don't be surprised if sometimes you see them referred to differently.



Some professionals, texts, and frameworks also include *recovery controls* as a control type. These controls could include safeguards such as backups, redundant systems, and so on. While you likely won't see this control type called out on the CRISC exam, simply be aware that sometimes you will see other control types, such as operational controls, in other contexts.

Figure 3.2 depicts the multidimensional nature of control categorization. As shown, a control can be preventive, automatic, and technical (e.g., a door with biometric access control).

FIGURE 3.2 Control characteristics and their implementation.



Control Standards and Frameworks

Control standards and frameworks are defined sets of controls that can be used in each set of circumstances, market, or area. Controls are often tied to governance directives and regulations, but this isn't always the case. Sometimes, different regulations or even process frameworks reference controls only to the extent that a control framework is required, without specifying which control standard must be used. The Health Insurance Portability and Accountability Act (HIPAA) is a good example of this scenario because it requires security and risk controls, but does not specify which control set must be used (although HIPAA guidance does mention the NIST control catalog). As another example, while references to controls are specifically called out in the ISACA Risk IT Framework (RR 2.1, "Inventory Controls," and RR 2.4, "Implement Controls"), the framework itself does not give any clear guidance on how controls are defined, selected, or implemented; however, the assumption is that implementing the framework will show you how to use and evaluate whichever control set you choose. The NIST Risk Management Framework (RMF), in contrast, does guide how its controls are constructed, used, and assessed, devoting several publications to this subject.

Specific industry or market environments may require context-specific controls, as indicated in Table 3.2.

Controls in each framework are often organized by area or context, usually by functionality or relevance. For instance, the CIS and NIST control standards (discussed shortly) are broken down into areas, or *families*, that apply to specific contexts. For example, the NIST Physical and Environmental Protection (PE) controls address physical and environmental concerns, such as physical access controls and environmental factors (lighting, humidity, and temperature). In contrast, the NIST Program Management (PM) control family addresses management processes that often appear as administrative controls, such as formally appointing a responsible information security officer. Most other control frameworks discussed in this chapter similarly categorize controls, usually by process area or other relevant aspect of context.

TABLE 3.2 Examples of areas requiring control standards.

Area or Environment	Corresponding Control Set(s)
U.S. federal government	FISMA, which prompted the creation of NIST Special Publication 800-53, revision 5
U.S. federal government related	FedRAMP, which is based on NIST Special Publication 800-53 CMMC, which is based on NIST Special Publication 800-171
Aggregated healthcare	HIPAA, HITECH, or MARS-E
Healthcare	HIPAA or HITECH, although HITRUST is sufficient from a legal perspective
Finance	FFIEC, SWIFT, NCUA, or relevant international controls
Credit cards	PCI DSS
Publicly traded	Sarbanes-Oxley (SOX)
Private/commercial	Center for Internet Security (CIS) controls, NIST CSF, ISO/IEC 27001, or AICPA's SOC2



You'll find that most control libraries similarly group controls into specific categories based on function or context, and most use specific identifiers for each control.

Which controls from the various frameworks an organization uses depends on several factors, including governance, market segment or area, level of regulation, and internal factors such as risk tolerance, risk appetite, and maturity of risk management processes. Controls are usually developed and applied to specific risk scenarios. These risk scenarios are developed by the organization or industry as part of their risk identification process. They describe threats, threat actors, events, time, and other circumstances, but, of course, control assignment and use may be required by legal or regulatory governance, as well as by the organization that intends to employ them. It's also common for an organization to use more than one control set or framework, as they may apply differently to various aspects of the business or mission areas. For example, a hospital may use NIST controls for some areas but may also be required to use PCI DSS controls since it may process payments from credit cards as patients pay for medical services. Often, controls from multiple sources are applied to one or more areas of the business.

In examining the various control frameworks, you'll also find that there are many similarities between controls in some instances; often, a control in one framework has a close equivalent in another framework. For example, CIS Control 10, "Malware Defenses," loosely maps to the corresponding NIST control, SI-3, "Malicious Code Protection," but describes a more generic process than the equivalent NIST control. There will also be controls that lack a matching control or context in another framework; these are often controls that address specific instances of risk or security requirements unique to a particular environment.



Most control catalogs include controls that can be easily mapped to controls in other catalogs. In some cases, this mapping may be imprecise, less detailed, or even map to multiple controls in another control library.

Now that you've learned about some generalities about controls, it's time to look at various control frameworks you might use in your professional career. Keep in mind that some control frameworks go into more depth than others; some controls in a given framework may treat a particular subject area lightly, whereas others may go into detail. Also, keep in mind that some control frameworks are geared toward specific environments or areas, and not all frameworks are necessarily geared toward security specifically; some frameworks apply to a broad general business or risk management area but may have security controls built into them for specific needs or direct security actions from a higher level. Not all control frameworks are technical in nature; some address security and risk from a business-process perspective rather than a technical one. Some control frameworks are privacy-oriented and provide a different set of controls. Some of the more popular control frameworks you may encounter in your professional career as a risk practitioner include NIST, PCI DSS, CIS Controls, and ISO/IEC 27001/27002.

NIST

The U.S. Department of Commerce's National Institute of Standards and Technology (NIST) has produced several documents—some of which I have already mentioned—that relate to risk management. The entire Risk Management Framework (RMF) and its components span several documents. I discuss a particular piece of the framework: the controls catalog. NIST Special Publication (SP) 800-53, revision 5, "Security and Privacy Controls for Information Systems and Organizations," is the control catalog supporting NIST's RMF. The NIST controls are a comprehensive set of security measures and processes intended to mitigate and reduce IT security risks. The NIST control catalog is divided into 20 separate control groups, called *families*, each related to a particular aspect of IT security.

The NIST control families span a great number of subject areas. Each family contains several controls; some families (such as the Awareness and Training family) may have as few as six controls, while others have many controls that go into great detail on security measures. An example is the System and Communications Protection (SC) family, with 51 controls.



You will likely not be tested on any specific NIST controls; however, it's a good idea to have this knowledge to differentiate among the different control frameworks you may see on the exam, as well as for real life.

PCI DSS

The *Payment Card Industry Data Security Standard (PCI DSS)* is a set of security requirements levied on merchants that process credit card transactions. The standard was developed jointly by major players in the payment card industry, including Discover, Visa, MasterCard, and American Express. The standard was developed to impose security requirements and controls on retailers (merchants and service providers) to reduce credit card fraud and identity theft. Note that PCI DSS is not a law; rather, it is an industry effort to self-regulate. Although not a law, it has the effect of being mandatory for governance purposes; retailers must comply with PCI DSS to process credit cards for major industry brands. A few U.S. states, however, have codified PCI DSS standards into their laws.

The PCI DSS (currently at version 4.0.1 as of this writing) consists of six control objectives, broken down into 12 major requirements, which are further divided into sub-requirements. Table 3.3 lists the basic PCI DSS requirements.

TABLE 3.3 PCI DSS control objectives and requirements.

Control Objective	Requirement
Build and Maintain a Secure Network and Systems	1. Install and Maintain Network Security Controls. 2. Apply Secure Configurations to All System Components.
Protect Account Data	3. Protect Stored Account Data. 4. Protect Cardholder Data with Strong Cryptography During Transmission Over Open, Public Networks.
Maintain a Vulnerability Management Program	5. Protect All Systems and Networks from Malicious Software. 6. Develop and Maintain Secure Systems and Software.
Implement Strong Access Control Measures	7. Restrict Access to System Components and Cardholder Data by Business Need to Know. 8. Identify Users and Authenticate Access to System Components. 9. Restrict Physical Access to Cardholder Data.
Regularly Monitor and Test Networks	10. Log and Monitor All Access to System Components and Cardholder Data. 11. Test Security of Systems and Networks Regularly.
Maintain an Information Security Policy	12. Support Information Security with Organizational Policies and Programs.

On the surface, these objectives and requirements don't seem to go into much depth. Understand, however, that PCI DSS further decomposes the requirements into detailed processes and procedures. The standard also guides testing at the control level. For example, requirement 5, "Protect All Systems and Networks from Malicious Software," breaks down into four sub-requirements, some of which are further subdivided. Requirement 5.2.1, continuing the example, specifies continuous monitoring and evaluation of hosts for malware threats. Overall, there are more than 200 sub-requirements, which are equivalent to controls.



You likely won't be asked any PCI DSS control–specific questions on the exam. Still, you should keep in mind its applicability to merchants and retailers that process credit card transactions. You can find more detailed information on the PCI DSS standards at the website of the PCI Security Standards Council at `https://www.pcisecuritystandards.org/`.

CIS Controls

The *Center for Internet Security (CIS)* publishes 18 critical security controls, which are essentially areas or families that decompose into 153 individual safeguards, spanning three implementation groups. These 18 areas are similar to other control sets in that they address common security concerns, including malware, penetration testing, software asset inventory, account management, access control, and so on. The three Implementation Groups (IGs) are designed to provide a path for progressively implementing safeguards, with prioritization and scalability in mind, to help smaller organizations, or those with limited resources, implement at least 56 core security controls that are necessary to secure systems and data.

The CIS controls were formally known as the SANS Top 20 Critical Security Controls, before the CIS took responsibility for them in 2015. They were originally developed by an effort led by the National Security Agency (NSA) and many other government and industry partners. The most current version of the CIS controls is version 8.1, published in June 2024.

ISO/IEC 27001/27002

Although I've primarily discussed different control sets originating in the United States (such as PCI DSS and NIST), most of the control frameworks map to, and in some cases originate from, standards used worldwide and developed by internationally respected governing bodies. In Chapter 1, I briefly mentioned some of these standards, including those from ISO and IEC. The ISO and IEC produce international information technology standards, including several that apply to information security and risk management. You may see individual standards prepared by each of these organizations, as well as combined standards, usually indicated as ISO/IEC for those that are jointly developed. I discussed standards that applied more specifically to risk assessments in Chapter 1, but there are also ISO/IEC standards that are essentially control requirements. Two examples are ISO/IEC 27001:2022, "Information security, cybersecurity and privacy protection—Information security management systems—Requirements," and its companion guide, ISO/IEC 27002:2022, "Information security, cybersecurity and privacy protection—Information security controls." These two standards (both updated in 2022), like those discussed previously, break down

their controls into categories, including information security policies, human resource security, asset management, access control, cryptography, and so on. Since the ISO/IEC 27001 and 27002 standards are developed by an international standards body, they are used worldwide.

As mentioned earlier, most of the control standards I have discussed include built-in mappings to facilitate their implementation, even though organizations may have multiple governance requirements to which they must adhere. Both the NIST and CIS controls have mappings to the ISO/IEC standards, for example. In examining these control mappings, a couple of points should stand out. First, there's not necessarily a straight one-to-one mapping of controls. In some cases, the requirements for a given control in a control set may map to one, many, or even none of the other controls in the control set. Second, controls tend to follow similar patterns and requirements, such as acceptable use, password requirements, media and storage, encryption, and so on. Most controls from the various control sets have these commonalities because all of them in some way trace back to the three goals of security (confidentiality, integrity, and availability) as well as the supporting elements of identification, authentication, and authorization. This is generally true regardless of which security control framework you examine.



You'll find that many control frameworks are derived from and map to the ISO/IEC standards. They are international standards and can be applied across most industries and market spaces.

Control Design, Selection, and Analysis

The importance of designing, selecting, maintaining, and monitoring the right controls cannot be understated. Remember that controls serve three primary purposes: protecting assets, ensuring compliance, and reducing risk. If controls are not adequately planned for and implemented, they will fail to meet any or all three of these purposes. In the next few sections, I discuss the critical elements of control rationalization, design, selection, and analysis.

Control Rationalization

Control rationalization is the systematic process of evaluating, consolidating, and optimizing an organization's control environment to ensure that controls are effective, efficient, non-redundant, and aligned with business objectives and governance requirements. As organizations adopt multiple frameworks, expand their technology stacks, or implement new security tools, their control environments often grow organically, resulting in duplicate controls, overlapping responsibilities, and unnecessary operational burden. Rationalization addresses these issues through structured analysis and decision-making.

Effective control rationalization supports several outcomes: reduced operational cost, decreased control fatigue, clearer ownership, improved auditability, and stronger alignment between controls and enterprise risk. For risk practitioners, rationalization also ensures that controls are appropriately mapped to risk scenarios and that effort is focused on those that materially reduce risk.

Step 1. Inventory and Mapping

The rationalization process begins by cataloging all existing controls across administrative, technical, and physical categories. Each control must be mapped to:

- Applicable risk scenarios
- Regulatory or contractual requirements
- Internal policies and standards
- External frameworks (e.g., NIST SP 800-53, ISO/IEC 27001, PCI DSS, CIS Controls)

Mapping identifies where controls overlap, where gaps exist, and where multiple controls address the same requirement. This step also reveals control areas that lack formal requirements, which may indicate misalignment with governance expectations.

Step 2. Assessment of Control Value and Effectiveness

Each control is evaluated based on its effectiveness and contribution to risk reduction. Assessment factors include:

- Strength of risk mitigation provided
- Actual versus expected performance
- Integration with other controls
- Dependence on external systems or processes
- Historical evidence of success or failure
- Cost to implement and maintain

Controls that provide limited risk reduction, demonstrate low maturity, or are difficult to maintain may be candidates for consolidation or removal.

Step 3. Identification of Redundant or Overlapping Controls

Rationalization seeks to reduce unnecessary duplication. Common redundancies include:

- Multiple technical controls performing identical filtering or monitoring
- Repeated policy requirements enforced through separate procedural controls
- Legacy controls still operating after new systems made them obsolete
- Project-specific or unit-specific controls that duplicate enterprise-level controls

Redundant controls increase operational workload, complicate audits, and consume budget without improving security.

Step 4. Consolidation and Optimization

Where redundancies exist, the organization may:

- Consolidate multiple controls into a single, enterprise-wide control
- Replace overlapping controls with a more effective or cost-efficient alternative
- Integrate controls into existing systems to reduce overhead
- Standardize configurations across business units for consistency

Consolidation decisions must consider interoperability, cost, compliance obligations, and business-unit dependencies. Any proposed control elimination must be validated to ensure that risk remains within appetite and tolerance.

Step 5. Assignment of Clear Ownership

Rationalization frequently uncovers ambiguous or duplicated ownership assignments. As controls are consolidated or redesigned, updated ownership must be established for:

- Risk owners
- Control owners
- Supporting process owners

This reduces confusion and ensures accountability throughout the lifecycle.

Step 6. Documentation and Governance Update

Updated control descriptions, mappings, testing procedures, and responsibilities must be documented in:

- The control catalog
- Policy and standards documents
- GRC system repositories
- Change management records

These updates ensure traceability and support future audits and assessments.

Step 7. Continuous Review

Control rationalization is not a one-time activity. Ongoing technology changes, regulatory updates, and business transformations require periodic reevaluation. Many organizations integrate rationalization into annual risk assessments or into major system updates to ensure the control environment remains efficient and aligned with enterprise risk.

Control Design

Controls are designed and applied to a system, data, process, or even organizational element to address a weakness or vulnerability or to counter a specific threat. Often, controls are designed to address a specific risk response strategy, such as risk reduction/mitigation, sharing, avoidance, and acceptance. Controls are applied to reduce the impact and likelihood of risk, as follows:

- Reduce the likelihood that a threat agent will initiate a threat (deterrence, preventive, and detection controls are examples).
- Reduce the likelihood of a threat exploiting a vulnerability (preventative and detection controls).
- Reduce the impact on the organization or asset if a threat exploits a vulnerability (corrective, compensating, and recovery controls).
- Reduce vulnerability (deterrent, preventive, corrective, and compensating controls all perform this action).

Note that none of the scenarios listed earlier describe a case where a control prevents a threat or a threat actor. Threats and threat actors exist, and controls can't really prevent their existence. Controls, however, can affect how they interact with vulnerabilities and assets. Also note that all these strategies may be pursued simultaneously when controls are well designed and effective, leveraging economies of scale and the multiple use of controls.

Control design should also reflect the principle of *economy of use*, ensuring that controls are straightforward, efficient, and minimally intrusive to legitimate workflows. Controls that impose excessive steps, require unnecessary user effort, or complicate routine tasks often lead to reduced adoption or unintended workarounds. Incorporating economy of use into design decisions helps ensure that controls remain effective in practice and sustainable over time.

A practical element of control design is establishing *control baselines*, which define the minimum set of controls required for a given system category, environment, or data sensitivity level. Baselines ensure that all systems start with a consistent, governance-aligned level of protection before any system-specific tailoring. By defining expected default controls—such as authentication requirements, logging standards, encryption thresholds, or network segmentation rules—baselines provide a reference point for evaluating design adequacy, identifying deviations, and determining whether additional or compensating controls are needed. They also streamline control selection by providing risk practitioners with a starting framework that aligns with regulatory requirements, organizational policies, and industry best practices.



Remember that controls can affect the likelihood or impact risk elements only by reducing vulnerabilities or protecting assets. They cannot eliminate or reduce threats or threat actors; they only reduce the likelihood or impact of an attempted or successful exploitation of a vulnerability by a threat.

Control Selection

When an organization selects controls to protect systems and minimize risk, several considerations arise. First, the organization should review its governance requirements and ensure its controls meet them, particularly where specific data protection requirements are specified. This may include the level of performance or functionality of a control (such as encryption strength). The organization also must look at its existing infrastructure. It is likely to already have some controls in place, although they might not be applied directly to the system in question. Take, for example, an organization that has industrial control systems and other operational technologies. Many of these devices cannot support basic security measures, such as authentication or encryption; however, as compensating controls, the organization could implement network segmentation by putting them in their own physically or logically separated subnets and install authentication and encryption mechanisms specifically for those devices, since they do not have their own built-in capabilities. If there are existing controls that already meet the requirements for a control, they should be used to the maximum extent possible.

Second, controls should also be selected based on their effectiveness, as indicated by the design parameters I discussed earlier in the "Control Design" section. A control that is minimally effective, and perhaps meets governance requirements, may still not be desirable if it does not provide the level of security the organization needs. Effectiveness may also mean how well the control reduces risk or ensures compliance. If either of these levels of functionality or performance is not met by the control, the organization should consider strengthening or changing the control.

Third, cost is also an issue with control selection. A control that costs $10,000 to implement but only protects a single system valued at $2,000 is not cost-effective. However, a control that costs $10,000 and protects several systems, even if individually they are not high-value systems, but collectively generate $100,000 in revenue for the organization per year, is likely more cost-effective. Control selection regarding cost should include both the control's cost and the cost to maintain it. This includes spare parts and components, service contracts, and the cost of qualified personnel to maintain the control. This is where quantitative data elements, such as asset value (AV), exposure factor (EF), single loss expectancy (SLE), and annualized loss expectancy (ALE), which you learned about in earlier chapters, come into play, since they affect whether a control is worth it financially to implement.

An important consideration during control selection is the organization's capability to operate and sustain the control. This includes the availability of skilled personnel, supporting processes, required technology platforms, and the maturity of existing operational practices. A control that appears effective on paper may fail in practice if the organization cannot implement or maintain it at the required level. Evaluating capability alongside cost, effectiveness, and compliance requirements ensures that selected controls can be executed reliably over time.

Ultimately, security and risk practitioners recommend controls based on their effectiveness and cost to management, who must decide whether to implement them. If the desired controls cannot be implemented for any reason, both practitioners and management must consider selecting appropriate compensating controls, which must be approved and documented.

Control Analysis

Controls are analyzed at several different points in their lifecycle. They are reviewed to ensure they fulfill their basic functions of asset protection, governance compliance, and risk reduction. They're examined not only for effectiveness (how well they do their job) but also for efficiency (how well they do their job with as little additional effort, resources, and management as possible). Even if a control is performing its job functions well, if the cost to maintain it is prohibitive, the organization may need to consider changing the control or its implementation.

The most common way to analyze controls is through control assessment, in which controls are evaluated for their effectiveness. A control can be evaluated using various methods, such as interviewing key personnel responsible for managing and maintaining the control, reviewing associated documentation, observing the control in action, and testing it through technical means. These different assessment methods will give a risk practitioner an idea of how effective the control is in meeting its primary functions.

Risk assessments for the specific system under review and for the entire organization will also indicate how well a control performs its function. Even if a risk practitioner is not examining a particular control, levels of increased or decreased risk for the system itself or the entire organization would indicate whether one or more controls are effective. Historical and trend analysis methods tell a risk practitioner how risk has changed over time. They also help the practitioner predict how risk will change in response to changes in the environment, technology, the threat landscape, and the effectiveness of controls. If risk has increased inexplicably, the practitioner must analyze the controls designed and implemented to mitigate it.

After controls have been tested and analyzed, the results of those analyses should be recorded in the overall risk register, particularly about the risk the control is intended to mitigate. Obviously, management must be informed of any changes in risk caused by control effectiveness issues.

Control Implementation

Controls aren't simply installed in infrastructures, even if there is an urgency to do so. As discussed previously, some planning is required for control implementation. Most of this planning focuses on determining the effectiveness of the control and ensuring it is designed to protect assets, assist in compliance, and reduce risk. There are also other considerations when implementing controls, including documenting and publishing controls, the change management process, configuration management, and testing.

Documenting and Publishing Controls

When an organization has finalized the controls to be implemented, it will publish those controls and announce the availability of the control documentation to affected personnel.

Control documentation generally includes the following information about each control:

- **Control name:** The name of the control
- **Control owner:** The position responsible for operating or supervising the operation of the control
- **Control language:** The actual control statement
- **Control description:** A description of the control, as the control statement may have ambiguities
- **Control objective:** The control objective that applies to this and possibly other controls
- **Background:** Other information about the control, such as why it was enacted
- **Control testing procedure:** Either the full testing procedure, or a reference to a separate testing procedure document
- **Control testing frequency:** How often the control should be tested
- **Applicability/scope:** The systems, processes, locations, and so on, where the control is required to be implemented

On the topic of scope, it is common for organizations to implement a set of controls applicable across all contexts. Such controls are known as *common controls* or *IT general controls (ITGCs)*.

Like other business documentation, such as processes and procedures, organizations conduct formal, periodic reviews of their controls to ensure they remain applicable and that the information remains current.

Change and Configuration Management

Change and configuration management processes are critical during control implementation because, in practice, implementing new controls involves infrastructure changes and must therefore be part of the formal change control process. This involves requesting the change and presenting a justification for it. In this case, the change is necessary to address a gap or need that an existing control (or the lack of a control) does not address. Compliance may also drive the need for change, as governance may require the control to be implemented. Risk reduction is also a valid justification for implementing a control in the change management process. Once the change is requested as part of the formal process, it is reviewed by the change control board and approved or returned for amendments to the request. After approval, it must go through the testing process. Once it clears the implementation testing process, discussed in the next section, the change is formally approved and implemented. If, for some reason, it does not pass the different types of tests it must go through before implementation, the change is returned to the requester so that any test failures can be addressed.

Control implementations may use different strategies, such as parallel installation and implementation with similar systems, gradual updates to the system, phased implementation, or even a clean cutover to the system using the new controls. Since the installation may incur unacceptable downtime for critical systems, this must be carefully considered. The organization must decide the appropriate strategy, based on its resiliency, system and data criticality, and tolerance for any risks that always come with implementing a significant change in the infrastructure.

Implementation Testing

Although controls undergo testing at various points in their lifecycle, testing before implementation is critical because it is often difficult to return to a previous control or a point in time before it is implemented. This is especially true if the infrastructure is significantly altered to accommodate the new control. Consider a network infrastructure that has a new firewall or other network security device installed. The network may have to be rerouted and reconfigured, and new rules configured on the network security device. If, for some reason, it is installed without testing and then fails, it may cause a serious impact on business operations if users cannot use the network, get on the Internet, or otherwise perform their duties.

Several types of implementation testing should occur on new controls before they are implemented. One of the first is *interoperability testing*. Interoperability is critical since the new control must work with existing technologies and infrastructure. If, for some reason, they do not work together, this can cause work stoppages or outages and will impact business operations. A risk analysis conducted before implementing a new control should take this into account, and the organization should be reasonably informed about how the new technology or control will interact with the existing infrastructure before it is tested and implemented.

Another type of testing involves evaluating security mechanisms inherent to the new control, such as authentication and encryption. Not only should they be interoperable with existing infrastructure, but they should also be tested for functionality and performance, as well as for compliance with any governance requirements. It would not do to install a security control that has legacy authentication or encryption mechanisms that do not meet governance requirements, for example, and therefore fails to protect systems and data.

Testing a control, especially a complex one, before implementation ensures it is interoperable, meets technical and governance requirements, and performs and functions as it should. Understanding the limitations and operating parameters of a complex control through testing before implementation is a method of risk reduction.

Other Implementation Considerations

It goes without saying that implementing a complex control involves many more aspects than are discussed here. Some of these are resource considerations, such as ensuring there is sufficient budget and enough qualified people to install, manage, and monitor the control, not to mention political considerations within the organization. Consider implementing a relatively "easy" control, such as revoking administrative rights on user workstations. Not only will users typically revolt, but managers accustomed to having administrative rights on their own workstations will complain and likely go to upper management, which may cave in to those complaints, effectively rendering the control useless. Users also typically balk at having their systems offline for any period of time, as it may interrupt their work. All these issues should be carefully considered and decided upon by management, or at the very least kept in mind by the risk practitioner.

Another aspect of control implementation is documentation. All steps of the implementation should be thoroughly documented, as well as any exception items, such as configuration or interoperability issues that may need to be addressed later. Documentation also includes policies and procedures to support the control, change management records, implementation test results, operator training materials, maintenance schedules, and any other information relevant to the control.

Control Testing and Effectiveness Evaluation

Control testing and evaluation, as well as analysis, were discussed earlier in the chapter and are worth repeating here. Controls must be tested and evaluated not only before implementation, but also at various points throughout their lifecycle. Controls should be tested and assessed as part of routine risk assessments, controls testing, vulnerability and penetration testing, and so on. Controls are tested not only for their effectiveness in protecting assets, but also for their compliance with governance and their ability to reduce risk to acceptable levels. Remember that earlier I discussed four methods for testing a control:

- Interviews with key personnel, such as operators, security analysts, engineers, and even managers
- Observation of the control in operation to ensure it does what it is supposed to do

- Review of any documentation supporting the control, such as policies, procedures, maintenance records, architectural diagrams, exceptions, and so on
- Technical testing, such as that conducted during vulnerability assessments and penetration testing

All of these different types of tests are critical for getting a complete picture of how well a control is functioning. The results of control testing should be documented and reported on in various circumstances as part of different assessments and tests.

Controls that fail testing or do not perform or function as expected should be further evaluated in more detail for any possible changes and any additional increase in risk. A control may need to be reinstalled, replaced, or upgraded, or compensating controls may need to be added to strengthen the function the control is meant to perform.

Continuous Monitoring and Assurance

Continuous monitoring and assurance are key components of an effective risk management program. While periodic assessments and control tests provide point-in-time evidence of control performance, continuous monitoring ensures that risks, controls, and the operating environment are evaluated on an ongoing basis. This allows the organization to quickly identify changes in risk conditions, detect control degradation, validate the effectiveness of risk treatments, and maintain alignment with governance requirements.

Continuous monitoring is both a technical and an organizational process. It requires a combination of automated data collection, periodic manual review, defined procedures, and clearly assigned responsibilities. Together, these mechanisms provide ongoing assurance that risk remains within appetite and that implemented controls continue to perform as intended.

Objectives of Continuous Monitoring

Continuous monitoring supports several core objectives:

- **Detect changes in risk posture** by identifying new vulnerabilities, emerging threats, or control deviations
- **Validate control performance** between scheduled assessments
- **Provide early warning indicators** of control failure or increased risk exposure
- **Maintain compliance** with regulatory and contractual requirements that mandate ongoing oversight
- **Support timely decision-making** by supplying current data on risk and control effectiveness

Monitoring assures that risk treatments remain effective over time and that the organization can detect and respond to changes in the environment.

Components of a Continuous Monitoring Program

A robust continuous monitoring program includes the following components:

- **Automated data collection:** Systems collect real-time or near-real-time data from logs, sensors, vulnerability scanners, configuration management tools, identity platforms, and network monitoring solutions. Automated collection ensures accuracy and reduces manual effort.

- **Control performance metrics:** Key performance indicators (KPIs) and key control indicators (KCIs) track ongoing control behavior. These metrics detect control drift, unexpected failures, and threshold breaches.
- **Risk indicators:** Key risk indicators (KRIs) provide insight into changes in the environment that may indicate increased likelihood or impact of a risk scenario.
- **Baseline and deviation analysis:** Controls and systems must have defined baselines (e.g., configuration settings, patch levels, access entitlements). Continuous monitoring compares observed conditions to approved baselines and flags deviations.
- **Event correlation and analysis:** Monitoring tools correlate events across multiple systems to detect patterns that may indicate emerging risk or control weaknesses.
- **Periodic manual review:** Some controls, such as business process controls, require periodic manual evaluation through reviews, attestations, or sampling.

Assurance Activities

Continuous assurance validates that controls remain effective and risk treatments continue to operate as intended. Assurance activities include:

- **Ongoing control testing** using automated or manual methods
- **Configuration and change monitoring** to detect unauthorized or risky modifications
- **Continuous vulnerability assessment** to identify newly discovered weaknesses
- **Threat intelligence integration** to detect relevant external risks
- **Policy and control attestation processes,** where control owners confirm continued operation
- **Independent review or audit** where required by governance or regulation

Assurance supports accountability by providing continuous evidence that controls operate effectively.

Integration with Risk Treatment and Decision-making

Continuous monitoring must be aligned with the organization's risk treatment plans and governance processes. Data from monitoring activities should:

- Identify when risk exceeds defined thresholds
- Trigger escalation procedures for risk owners
- Inform adjustments to existing risk treatments or controls
- Support formal review and revision of risk scenarios
- Feed into dashboards, scorecards, and other reporting mechanisms

Monitoring ensures that management receives accurate, up-to-date data needed for timely, risk-based decisions.

Relationship to Other Frameworks

Many frameworks incorporate continuous monitoring, including:

- **NIST Risk Management Framework (RMF)** (Step 7: Monitor)
- **NIST SP 800-137** (Information Security Continuous Monitoring)
- **COBIT** (Managed Assurance, Managed Risk, and continuous control monitoring objectives)

Alignment with these frameworks reinforces consistency, compliance, and audit readiness.

Continuous Improvement

Continuous monitoring is not a static activity. As technology, threats, or business processes evolve, monitoring scope, metrics, and automation capabilities must be updated. Continuous improvement ensures that the monitoring program remains effective and adapts to emerging risks and changes in the control environment.

Risk Monitoring and Reporting

What are the risk reporting requirements for your organization? Are they dictated by government regulations, international business partnerships, or other binding agreements? Understanding what is required of your organization in terms of *risk monitoring* and reporting is critical to fulfilling governance requirements, whether imposed by external sources, such as laws and regulations, or by internal governance, such as policies issued by executive management. You also need to understand what you should report in any of these contexts. For example, it is a good idea to report any ongoing risk issues, the status of any ongoing remediation actions, how your controls are performing, any incidents, and how your overall risk response activities are performing.

Risk Treatment Plans

While reporting may at first seem to be only an occasional aspect of risk management, it takes place on a continual basis. Obviously, risk assessments, control assessments, and different types of tests have their own reporting requirements that will be accomplished when those test events are completed. However, reporting can also take place through continuous information sent to management.

Some of the most critical pieces of information to report to management, regardless of a point-in-time test event, include risk response and treatment options, and the plans used to implement those options. Management should always have the opportunity to review risk treatment options as they are developed, and risk practitioners have a responsibility to formally present those recommendations to management so that leadership can make appropriate decisions regarding those treatment options. Understand that simply because

a risk practitioner recommends specific treatments or responses does not necessarily mean management will approve those options. Alternative risk treatment responses may be approved instead, based on information the risk practitioner may not have or on management's priorities. This is particularly true when risk treatment options are expensive, and management is seeking to reduce costs while balancing effective risk response. It is also true that management may need to make a broader decision, such as risk avoidance, for example, by taking an entirely different strategic path within the business. When reporting on risk treatment options, you should always present the best recommendation and also develop alternatives for management to consider. In any event, it's important to convey risk treatment options and their progress to management on a continual basis, not merely after they have been implemented.

A risk treatment plan is a structured document that defines how an organization will address identified risks to achieve alignment with its risk appetite, tolerance, and overall governance requirements. While risk response options (mitigation, sharing, acceptance, and avoidance) describe what type of response will be pursued, the risk treatment plan specifies how the chosen response will be implemented, monitored, validated, and reported. The plan provides traceability, accountability, and a repeatable method for ensuring that risk responses are executed effectively and consistently across the enterprise.

A comprehensive risk treatment plan should include the following components.

- **Risk description and context:** Every treatment plan begins with a clear description of the risk being addressed, including the risk scenario, affected assets, threat actors, vulnerabilities, potential impacts, and the analytical results (qualitative or quantitative) that justify treatment. The scenario should be written in a structured, repeatable format; for example, describing the actor, action, asset, event, and business consequence. The context must also specify any regulatory, contractual, or internal governance requirements that influence how the risk must be treated.
- **Risk owner:** The risk owner is accountable for ensuring the risk is treated appropriately. The plan must identify the individual or role responsible for this. Risk owners approve treatment decisions, allocate resources, and ensure that activities remain consistent with business objectives. Clearly identifying the risk owner prevents ambiguity and strengthens governance and accountability.
- **Selected risk response and rationale:** The risk treatment plan should identify the selected risk treatment option (mitigate, transfer, accept, or avoid) and summarize why that option was chosen. The rationale may include factors such as cost-benefit analysis, technical feasibility, compliance obligations, operational constraints, or alignment with risk appetite. Where multiple treatment options are combined (e.g., partial mitigation and partial sharing), the plan must document how each option applies to different aspects of the risk.
- **Required controls or actions:** If the selected treatment involves mitigation, the plan must list the tasks to be performed, and the controls to be implemented, strengthened, or retired. Each control should be described with enough detail to support testing and validation. For acceptance, the plan must document the residual risk that will remain, and the monitoring required to manage it. For risk avoidance, it should identify the

activity or condition that will be discontinued. For transfer, the plan should specify contractual obligations, service-level agreements, insurance coverage, or third-party arrangements.

- **Control owner(s):** Control owners are responsible for the design, implementation, operation, and maintenance of the controls associated with the risk treatment. Many risks require multiple controls spanning different organizational units; therefore, the plan must assign each control to a specific control owner. This distinction between risk owners and control owners ensures accountability for both decision-making and operational execution.
- **Implementation steps and milestones:** Each treatment action must include detailed implementation steps, timelines, and milestones. These should be realistic, measurable, and aligned with project management practices. Milestones help track progress and provide checkpoints for management oversight. Dependencies, such as technology procurement, workforce availability, or third-party activities, should be documented.
- **Resource requirements:** Treatment activities often require funding, staffing, tools, training, or external services. The plan must identify the resources necessary to implement the controls or actions, including estimated costs and any budgetary constraints. When resources are insufficient, the plan should document the shortfall so management can make an informed decision on prioritization or approval.
- **Success criteria and acceptance criteria:** The plan should define measurable criteria for treatment success. These may include technical performance parameters, reductions in likelihood or impact, compliance achievement, or improved control maturity. Acceptance criteria establish the point at which management agrees the treatment has been sufficiently completed. These criteria support transparency and ensure that treatment outcomes can be objectively validated.
- **Residual risk statement:** After treatment, some level of residual risk will remain. The plan must document the expected residual risk and demonstrate that it falls within the defined risk appetite and tolerance. If residual risk exceeds tolerance, the plan should identify additional actions, compensating controls, or escalation requirements.
- **Monitoring and reporting requirements:** Risk treatment does not end at implementation. The plan must specify how the risk and associated controls will be monitored. This includes KRIs, KCIs, KPIs, control testing frequency, periodic review intervals, and required reports. Monitoring ensures continued alignment with changing threats, technologies, and business conditions.
- **Validation and verification activities:** Validation determines whether the treatment achieved the intended risk reduction. Verification ensures that controls were implemented as designed and are functioning correctly. This may involve technical testing, audits, assessments, configuration reviews, or continuous monitoring. The plan should describe who performs these activities and how results are recorded.
- **Documentation and record retention:** All aspects of the treatment plan must be documented, including decisions, approvals, exceptions, changes, and results. This supports governance, auditability, regulatory compliance, and consistency in future assessments.

Data Collection, Aggregation, Analysis, and Validation

To make recommendations on risk treatment, risk practitioners must have a wide range of data at their disposal. The data must be trustworthy, accurate, and complete. A risk practitioner can leverage multiple data sources. Technical data, such as logs, can show how a control is operating. Real-time data regarding intrusion detection alerts, user behaviors, and so on is also helpful. Other data, including risk reports, incident reports, change and configuration management documentation, and information derived from control assessments, should be used to inform risk management decisions. The vast multitude of data available to the risk practitioner, however, might not be useful to an executive manager. It's the job of the risk practitioner to take all this data and distill it into useful, actionable information that clearly and concisely conveys how effective a control is, how it meets compliance requirements, and how it reduces risk. This information should also include details on changes in risk and how they affect the organization.

Regardless of the source of the data used to perform risk management activities, the data must be collected in a sound and verifiable manner and validated as accurate. This is where formalized and documented data collection and analysis processes can assist. Some data may be collected ad hoc, but for the most part, it should be collected regularly from verified sources and validated using various integrity-checking methods. The data must also completely describe the situation in the proper context and, of course, be analyzed by qualified risk practitioners and cybersecurity personnel. Data may be synthesized into useful information by reducing it, clarifying it, and performing trend or historical analysis. In any event, information supplied to management needs to present a clear, concise, and accurate picture.

Risk information is only as reliable as the data used to produce it. In many organizations, data feeding risk assessments, KRIs, KCIs, KPIs, and reporting is pulled from a wide collection of operational systems—identity platforms, logging environments, asset inventories, vulnerability scanners, GRC tools, financial systems, and third-party feeds. These data sources evolve independently, and their quality varies widely. For this reason, evaluating the quality of the data itself is an essential step to ensuring that risk reporting reflects real conditions rather than assumptions or partial visibility.

Several characteristics are commonly used to assess the quality of data within risk management programs. Although terminology varies among frameworks, the following dimensions represent the core attributes that practitioners rely on when determining whether a dataset can be used confidently:

- **Accuracy:** Data should reflect actual conditions, not approximations or outdated assumptions. A configuration database that reports a system as "encrypted" when the encryption module has been disabled, or a log source that attributes events to the wrong system, undermines the risk picture. Accuracy is especially important for data used in KRIs, because even small deviations can distort trend analysis or mask early signs of control drift.
- **Completeness:** Gaps in data coverage can be more damaging than inaccurate data because they often go unnoticed. Missing entries in access control lists, unscanned network segments, or partially collected log streams can leave entire risk scenarios

unmonitored. Completeness requires not only that individual records contain all required fields, but also that all relevant systems, applications, and processes are represented in the dataset.

- **Timeliness:** Risk conditions shift rapidly, and decisions made on stale information frequently lead to misaligned responses. A KRI based on vulnerability data that is several weeks old has limited predictive value. Timeliness does not imply real-time collection for all data types, but it does require that the data be current enough to support the sensitivity of the decisions being supported.
- **Consistency:** Data used for risk analysis must agree across systems and over time. Inconsistent identifiers, control status codes, severity definitions, or asset classifications introduce friction into analysis and can lead to contradictory conclusions. Consistency also matters when the same data element is used across multiple reports; a control shown as "implemented" in one dashboard and "not implemented" in another erodes confidence in both.
- **Integrity:** The reliability of data depends on its protection from unauthorized alteration or corruption. Integrity considerations often arise with log data, security telemetry, and control evidence. If the data can be manipulated—intentionally or otherwise—its use in audits, investigations, or reporting is compromised. Integrity measures include access controls, hashing, secure transmission, and validated data pipelines.
- **Lineage:** Understanding where data originates, how it is transformed, and where it ultimately resides is essential for determining its trustworthiness. Lineage reveals whether data passes through systems that alter its structure, introduce normalization rules, or aggregate values in ways that change meaning. When practitioners understand lineage, they can explain how a datapoint reached its current form and ensure that conclusions drawn from it remain defensible.
- **Validity:** Many datasets appear complete and current, but still fail basic validity tests. Validity concerns whether the data conforms to required formats, ranges, and rules. Risk systems commonly encounter invalid date stamps, malformed IP addresses, unclassified assets, or control evidence that fails to meet submission criteria. Invalid data may be excluded from automated pipelines, producing gaps that surface later as inconsistencies.
- **Uniqueness:** Duplicate records distort metrics, inflate incident counts, and skew performance indicators. Uniqueness ensures that each record corresponds to a single, distinct entity. This dimension is particularly important in asset inventories, incident databases, and identity repositories, where duplication can cause a cascading series of analytical errors.

Evaluating data against these dimensions is not a one-time exercise. As tools change, new sources are added, or collection methods evolve, data quality must be reassessed. Many organizations formalize this process through data quality scorecards, automated validation scripts, periodic stewardship reviews, or integration checks within their GRC platforms. Without such practices, even the most sophisticated risk monitoring and reporting functions can produce results that appear precise but are fundamentally unreliable.

Risk and Control Monitoring Techniques

The techniques used to monitor and assess risk, as well as control effectiveness, should be part of a formalized methodology adopted by the organization. This is where the overall risk strategy, along with its associated policies and procedures, comes into play. The risk strategy is management's commitment to risk management and provides an overall long-term direction for how the organization wants to manage its risk. The risk strategy is a high-level document and does not necessarily detail the day-to-day risk management activities an organization must engage in. The risk management policy, on the other hand, should provide more detail on the overall risk management program. This policy should be issued by management, in consultation with the risk practitioner. It should assign roles and responsibilities, as well as dictate the risk management, assessment, and analysis methodologies the organization officially uses. The policy can also define risk treatment options and how they should be considered. Most of the details required to implement the policy and, in turn, the strategy will typically be included in various lower-level risk management procedures, such as procedures used for risk assessment. In any case, the organization's risk management and monitoring processes should be standardized and approved by management for consistency, reliability, and predictability.

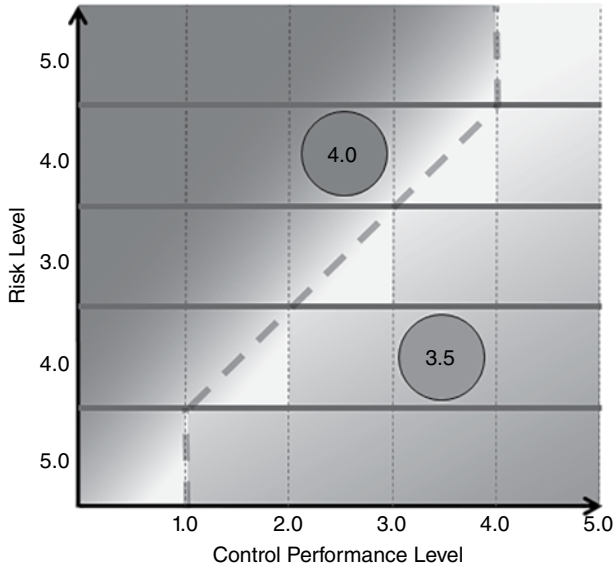
Risk and Control Reporting Techniques

Reporting on risk and controls may use various techniques to present data clearly, accurately, and completely. Some techniques lend themselves to better reporting than others; narrative reports are often undesirable except where a great deal of technical data is needed to support a recommendation or decision. Graphics and charts, on the other hand, are sometimes better suited to convey the bottom line to management without all the unnecessary details of a technical explanation. Other data presentation techniques, such as statistical, historical, and trend analysis, in the form of graphs or bar charts, may be useful for presenting data effectively so that management can make an informed decision.

Some of these data reporting techniques include heat maps, scorecards, centralized dashboards, and real-time information feeds. Heat maps focus on data that shows areas of red, yellow, and green; for example, data points in "red" (upper left) are the obvious areas of concern, and those in the "green" (lower right) areas are of lesser concern. Figure 3.3 shows an example heat map.

Scorecards are used to collect data and provide comparisons with other data, showing relative performance or effectiveness improvements compared to previously collected data. Centralized dashboards, along with real-time information feeds, aim to provide a single, one-stop view of the data points risk practitioners and managers need to monitor, often in real time, the status of control effectiveness, compliance, or risk. Often, however, these dashboards can seem to be very complex and "busy" with information, so extra effort must be taken to ensure they only display critical information, in the most simplistic manner possible, while allowing a risk practitioner or manager to "drill down" further to obtain more details about a piece of data in the dashboard.

FIGURE 3.3 An example of a heat map.



Regardless of the data presentation technique used, reporting must be clear, accurate, and complete, and it must convey the critical information needed for managers to make meaningful risk-based decisions.

Stakeholder-specific Reporting Needs

Along with understanding the organization's risk environment, it is important to understand your stakeholders and their reporting needs. For example, your board of directors will likely want a high-level, quarterly presentation on compliance levels relative to the previously discussed risk and control environment, as well as security performance relative to business goals and objectives. Business process owners, conversely, may want a much more detailed report on the previously identified individual risks, the new risks that could come to the forefront, and how different business units are supporting the risk management program across the entire organization. You must understand how these different stakeholders need different reporting mechanisms and tools, and clearly identify how they expect you to support them through timely, relevant risk reporting.

Standard Elements of Risk Reports

When reporting, you can include a few standard items, such as measured control effectiveness (which can be rolled up to a higher level, as necessary, for management), any ongoing issues, any perceived gaps, the status of any remediation actions, any incidents or risk events that have occurred since the last report, and the overall trending status of the

effectiveness of the risk management process. Also, identify ahead of time whether your report will be included as input to a larger enterprise report, and work to include any necessary information to facilitate that enterprise-level view. This will help garner support for the risk program supporting the business goals and objectives.

Incorporating Audit and Assessment Results

You should review the results of any third-party analysis or audit that has taken place and incorporate the findings into your understanding of the organization's risk posture. Has your organization conducted any type of internal audit or self-assessment? These will be important to review to see how they affect the risk profile and where they will need to be included. I've discussed the organization's risk tolerance in many places throughout this book, so suffice it to say that it should be considered a key input throughout the reporting process.

Creating Actionable and Decision-ready Reports

In the end, in order to develop a report that is valuable to management, a risk professional needs to make sure they have a clear understanding of the exposures their organization faces, the active mitigations for those exposures, and the indicators being used to measure the success or failure of the organization's implemented control set, as well as the current risk posture as measured against the risk appetite. All of this needs to be communicated to the stakeholders. The importance of articulating risk should not be underestimated; this is a critical task that ensures management receives accurate, unbiased information that is both actionable and aligned with the overall organizational risk strategy.

Responsible reporting involves helping leadership understand the risks to the organization and the criticality of reducing them by implementing appropriate controls, as well as monitoring and managing risk. This needs to be done in the clearest and most concise manner possible, understanding the concepts of risk, the indicators chosen, and the current statuses that enable leadership to make the best-informed decisions possible. Without the best possible information, we cannot expect organizational leadership to make informed decisions. You will also need to interpret technical findings and perhaps portray information in a manner that is less technical and more business-focused. By doing this, you will not only make the point that the risk program is useful, worthwhile, and supports business goals and objectives, but you will also identify opportunities for the risk program (and associated reporting) to integrate throughout the enterprise more cleanly.

Communicating Risk Scenarios and Business Impact

No communication is complete without thoroughly understanding and depicting both the worst-case and most probable scenarios involved with a risk event. A risk event could lead to many potential losses, not just those affecting the company's financial bottom line. There could be losses to the company's goodwill, including any associated brands. Think about major security breaches in previous years, such as Target, T.J. Maxx, and Sony. The result of their losses was both tangible (financial) and intangible (brands associated with a major

breach), which is difficult to quantify. This is also the time to include any other information pertinent to supporting the recommendations you've made, such as situations where a financial gain could be made (through exploiting an opportunity), or any outlying legal or regulatory concerns you might have.

Risk monitoring must also account for *risk realization*, the point at which a potential risk transitions into an actual event with measurable impact. When a risk is realized, the incident provides concrete evidence about the adequacy of existing controls, the accuracy of earlier assessments, and the effectiveness of risk treatment decisions. Realized risks should be captured in monitoring reports, analyzed for root causes, and incorporated into updated risk scenarios so that residual risk, KRIs, and treatment plans reflect the organization's true operating conditions.

A regional healthcare provider experienced a data exposure incident after several externally facing servers were found to be missing security patches. The organization had a documented vulnerability management process, but responsibility for patching was distributed across multiple business units with inconsistent workflows. During routine monitoring, a KRI showed a sustained increase in overdue patches, but no corrective actions were taken. Weeks later, attackers exploited an unpatched system, resulting in unauthorized access to protected health information.

Post-incident analysis revealed that the risk had been identified previously, but the organization lacked the capability and process maturity to enforce patch compliance across decentralized teams. The realized risk triggered a reassessment of ownership, consolidation of patch responsibilities, and implementation of a standardized control baseline. Continuous monitoring thresholds were also revised to ensure escalation when KRIs indicated elevated exposure.

Key Performance Indicators

I have discussed the necessity to test, assess, and evaluate controls as well as the risk itself. Remember that you're trying to ensure that controls are effective, meet compliance requirements, and reduce risk. To determine those things, you must measure various aspects of the control's performance characteristics, both positive and negative. For this, you use various *indicators* to help measure these items. Once you identify these indicators, you can move toward the indicators associated with measuring and reporting performance and current risk levels accurately and efficiently.

You will often hear indicators referred to as *metrics*, and in this text, those terms are interchangeable for our purposes. However, the types of indicators, and what they measure, differ. *Key performance indicators (KPIs)* are used to understand and measure control performance. Control performance can be anything from how often your firewall blocks malicious incoming traffic to how often your continuity of operations plan is exercised.

The level of "performance" can vary across organizations; you must determine performance goals unique to your organization, based on its performance goals, risk appetite, and tolerance levels.

So why do you need indicators? How do you know if you're doing well (or poorly)? Do you trust your risk team, hoping they have a reporting mechanism that will present stakeholders with the correct information? At the end of the day, you want to be able to monitor and analyze those indicators to track trends, assess the effectiveness and efficiency of the controls, and make decisions on creating new controls, ditching others, and sounding the alarm when risk has passed a pre-defined threshold.

Understanding what a KPI means to your organization requires asking a few questions. What is your organization's mission? Are you a research organization attempting to develop a new cryptographic algorithm? What things are important to your organization? Does it value availability above all other attributes (that is, the ability to conduct business through any fire, storm, hacking attempt, or any other scenario it may encounter)? Or does your organization value data privacy and security above any other consideration? These are questions you need to ask yourself when developing KPIs. Remember that these KPIs help you understand how controls are performing to enable all these attributes. You want to make sure the controls are performing as expected.

Ideally, KPIs should be "S.M.A.R.T" to be most effective. In other words, KPIs should be

- **S**pecific
- **M**easurable
- **A**ttainable
- **R**elevant
- **T**imely

When a KPI is S.M.A.R.T., it is tailored to the organization and the area being measured. The first question you need to answer is, what are you trying to measure? You should be, according to the S.M.A.R.T. methodology, *specific*. Are you attempting to measure availability? Then you may want to look at a metric that measures how often your systems go down—or, conversely, what is the percentage of time your systems are up?

A KPI also needs to be *measurable* in a way that ensures accuracy. If available, how would you measure it accurately? Perhaps you could determine how often the systems went down based on the `uptime` command in the UNIX operating system.

When a KPI is *attainable*, it is readily achievable. For example, you wouldn't want to set a performance goal of 100% uptime in a situation where you have an entirely new staff, new equipment, new management, and so on. Perhaps 85% is more achievable, and then you can work toward a higher goal.

KPIs that are *relevant* matter to your organization. If you're in a situation where time is of the essence, such as stock trading, then availability is essential. Without availability, your systems cannot buy or sell stocks, and your user base will quickly move on to another organization that values availability as much as they do. However, in other contexts, such as educational institutions, availability may not matter as much. Usability might be of

more concern, or the ability to access content via mobile applications or other cutting-edge technologies. Take the time to consider what's important or relevant to your organization when determining your KPIs.

Finally, *timely* KPIs allow data to be collected when needed, enabling you to quickly assess the performance of your measured initiative. You should be able to obtain data for the metric relatively easily from existing sources and repositories.

In some instances, a KPI may not meet all the criteria of S.M.A.R.T. For example, a firewall may block many external attacks per month. Depending on the organization, there may be nothing to attain, but most of the S.M.A.R.T. criteria are still met. This is perfectly fine in some circumstances.

Key Risk Indicators

Risk indicators are metrics, much like KPIs, used to determine whether and when the organization has a high probability of exceeding its defined risk appetite. Not all risk indicators, however, are key. *Key risk indicators (KRIs)* are highly probable indicators designed to accurately predict important levels of risk based on defined thresholds.

Along with determining the performance level of your controls, you need to identify ways to indicate the risk your organization is subject to. KRIs are related to, but separate from, KPIs, in that KRIs focus on how the metric predicts current and future risk rather than on the past performance of a specific control or set of controls. KRIs are linked to specific risks and provide an early warning that a risk is emerging, allowing leadership to be proactive and address it sooner rather than later.

ISACA's Risk IT Framework includes the following criteria to select KRIs:

- **Impact:** When a risk indicator has a high impact, it is likely to be a KRI.
- **Measurability:** When all else is equal, the easiest indicator to measure from is the best.
- **Reliability:** The indicator reliably predicts an outcome.
- **Sensitivity:** The indicator accurately indicates variations in risk.

Like the determination of KPIs discussed earlier, KRIs should be S.M.A.R.T., and they should drill down to the root cause and not just focus on the symptoms that are immediately evident. KRIs should include lagging indicators (which indicate current risk after an event has occurred), leading indicators (which indicate the controls in place to prevent an event from occurring), and trends (which consider patterns in the data that indicate risk). Root-cause analysis is also important here to determine the root cause of any risk, not just the symptoms. This is where risk assessments and analyses, control evaluations, and other data collection and testing come into play to provide you with this information.

Key risk indicators can be categorized as *leading* or *lagging* indicators. Both types are necessary to provide a complete view of risk behavior over time. Leading indicators help predict risk conditions before an adverse event, while lagging indicators measure the effectiveness of controls or the severity of risk after the event. Including both types enables proactive and retrospective analysis.

Leading Indicators

Leading KRIs measure conditions or behaviors that increase the likelihood of a risk event. These indicators provide early warning signals and allow management to intervene before impact is realized. Examples include:

- **Percentage of critical systems with overdue patches.** A rising value may indicate elevated exposure to exploitation.
- **User account provisioning backlog.** Delays in provisioning or deprovisioning may indicate increased identity-related risk.
- **Number of privileged-access exception requests per quarter.** Increasing trends may signal weakened access governance.
- **Volume of anomalous authentication attempts.** An upward trend suggests increased threat activity targeting identity systems.
- **Third-party security score degradation (e.g., from a continuous monitoring service).** Declining scores may indicate elevated supplier risk.

These indicators help identify when risk is approaching or exceeding tolerance thresholds, enabling earlier action.

Lagging Indicators

Lagging KRIs measure outcomes that have already occurred and indicate the realized level of risk or effectiveness of existing controls. Examples include:

- **Number of confirmed security incidents involving compromised user accounts.** Reflects the impact of prior weaknesses in authentication or monitoring.
- **Mean time to remediate high-severity vulnerabilities.** Measures the organization's responsiveness after risks have manifested.
- **Unplanned downtime hours attributed to control failures.** Indicates the effectiveness of operational resilience controls.
- **Number of regulatory non-compliance findings identified in audits.** Reflects compliance risk and the maturity of governance processes.
- **Financial loss attributed to fraud or cyber events within a reporting period.** Measures realized business impact.

Lagging indicators do not predict future conditions but are critical for validating whether the organization's risk management processes are working as intended.

Using Leading and Lagging Indicators Together

Both types of KRIs should be aligned with the same underlying risk scenario to ensure coverage across the risk lifecycle. For example:

- A *leading* KRI might track the number of unpatched critical vulnerabilities on external systems.

- A *lagging* KRI for the same scenario might measure the number of successful exploit attempts or security incidents involving those systems.

Using combined indicators allows organizations to monitor early warning signs, assess control effectiveness, and validate that risk treatments are achieving the required reduction in likelihood and impact.

Key Control Indicators

As discussed previously, KPIs are used to understand and measure control performance. *Key control indicators (KCIs)* are often used interchangeably with KPIs, but they are really a subset of KPIs, since some KPIs don't necessarily relate to controls. KCIs are those specific indicators that show the relative effectiveness of controls. This effectiveness, as mentioned, can be measured by how well a control protects an asset, how well it supports compliance with governance (e.g., the number of compliant versus non-compliant controls), and, probably most importantly in this context, how well it reduces or mitigates related risk. It is important that the control indicator focuses on the specific control and the particular risk it is meant to mitigate. Associating a KCI with a broad spectrum of risk is not effective; it should be associated with a very specific risk scenario. For example, a KCI could indicate the effectiveness of the control in terms of risk reduction by showing how it reduces the impact of a particular risk scenario for a given asset, the likelihood that a threat will exercise a particular vulnerability, or the vulnerability itself.

KCIs should be developed based on the control's known performance tolerances. You should determine the normal operating parameters for the control and what should be expected based on its configuration and implementation. Baselining control performance is important so that you understand what normal operation looks like. Abnormal or anomalous operation, such as a sudden spike in unknown network traffic, may indicate that the control is not operating within its normal range. This, in turn, shows the control may not be as effective as it should be.

Examples of key control indicators might include, but are not limited to, the following:

- Bandwidth utilization
- Number of false positives or false negatives detected
- Number of instances of malware detected
- Number of incidents per month attributable to user complacency or policy violations
- Number of vulnerabilities discovered during a scan

Additionally, not all KCIs need to be technical, although those likely make up most of the ones with which many organizations are concerned. KCIs can be developed for not only technical controls but also administrative and physical controls.

Summary

After risks are assessed and analyzed, organizations must choose appropriate responses. Risk mitigation, sharing, acceptance, and avoidance are the primary treatment options. Mitigation reduces risk through controls; sharing transfers some financial or operational exposure but not accountability; acceptance documents residual risk; and avoidance removes the activity that generates the risk. Risk and control ownership must be clearly assigned, particularly for common controls and ITGCs that support multiple business units.

Third-party risk management requires evaluating the threats, vulnerabilities, and obligations associated with external providers. Organizations must consider legal liability, data ownership, contractual protections, and the fact that third parties experience many of the same security and operational risks as internal systems.

AI and ML systems introduce unique risks seldom seen before. These risks include bias, hallucination, drift, and a lack of explainability. Training data must be carefully selected, particularly data within the scope of applicable privacy laws, to ensure continued compliance with those laws.

Issues, findings, and exceptions require structured handling to prevent unmanaged risk. Vulnerabilities that cannot be remediated, audit findings, or policy exceptions must be documented, risk-assessed, formally approved, and tracked through closure.

Risk is dynamic. Emerging technologies and changes in the operating environment can introduce new vulnerabilities or alter existing scenarios. Implementations should be based on business needs, and risk responses must adjust as conditions evolve.

Controls, including administrative, technical, and physical, reduce risk and support compliance. Their functions include deterrent, preventive, detective, corrective, and compensating actions. Controls are selected and designed based on effectiveness, cost, interoperability, and alignment with governance requirements. They must be implemented through disciplined change and configuration management processes, then tested initially and periodically to confirm continued effectiveness.

Ongoing risk monitoring and reporting ensure that management receives accurate, complete, and timely information. Reporting requirements stem from internal governance and external obligations. Risk practitioners must tailor reports to stakeholder needs, using dashboards, scorecards, or narrative formats. Key performance indicators measure process effectiveness, key risk indicators show changes in exposure, and key control indicators demonstrate whether controls are functioning as intended.

Exam Essentials

Understand risk treatment and response options. Organizations select from four primary risk response options: mitigate, share (transfer), accept, and avoid. Mitigation uses controls to reduce likelihood or impact. Sharing transfers some financial or operational exposure to a third party but does not transfer accountability. Acceptance is a formal decision to retain

residual risk within defined appetite and tolerance. Avoidance eliminates the activity that generates the risk.

Recognize the distinctions between risk owners and control owners. Risk owners are responsible and accountable for managing specific risks and approving risk treatments. Control owners operate and maintain controls that reduce the likelihood or impact of risks. These roles may differ, particularly when common controls span multiple business units.

Understand third-party risk management concepts. Third-party risk management includes identifying threats, vulnerabilities, contractual obligations, legal liabilities, and performance dependencies. Service-level agreements (SLAs), business associate agreements (BAAs), shared responsibility models, and periodic due diligence assessments (e.g., SOC reports, security questionnaires) are used to manage this risk.

Know how issues, findings, and exceptions are managed. Issues, audit findings, and exceptions represent deviations from required security baselines or controls. Effective management includes identification, triage, approval, documentation, tracking through closure, and periodic review. Exceptions must be formally approved, time-bound, and monitored.

Understand the management of emerging risk. Emerging risk results from changes in technology, business processes, the threat landscape, and the operating environment. New technologies introduce interoperability, compatibility, and security challenges that require reassessment of existing controls and risk treatments.

Know control types and control functions. Controls fall into three categories: administrative, technical, and physical. Common control functions include preventive, detective, corrective, deterrent, compensating, and recovery. Controls are selected and designed to reduce risk, ensure compliance, and protect assets. Understand that there is no single standard that depicts the multidimensionality of controls.

Understand major control frameworks. Common frameworks include NIST SP 800-53 Rev.5, ISO/IEC 27001:2022 and 27002:2022, CIS Controls v8, and PCI DSS 4.0.1. Frameworks define control requirements, support compliance, and enable standardized implementation across regulated sectors.

Comprehend control design, selection, and analysis. Controls are selected based on governance requirements, effectiveness, cost, interoperability, and risk reduction. Analysis includes evaluating implementation, performance, efficiency, and alignment with business objectives.

Understand control implementation, testing, and change management. Controls must be implemented through formal change and configuration management processes. Pre-implementation testing includes interoperability, security functionality, and performance. Post-implementation testing includes observation, documentation review, interviews, and technical assessment (e.g., vulnerability testing).

Recognize the requirements for risk monitoring and reporting. Monitoring includes ongoing evaluation of controls, KRIs, KCIs, and KPIs (key risk indicators, key control indicators, key performance indicators), using validated data sources. Reporting must be

accurate, complete, timely, and tailored to stakeholder needs through heat maps, dashboards, scorecards, or narrative summaries.

Understand risk treatment plans and their components. Risk treatment plans document assigned owners, required actions, resources, timelines, acceptance criteria, dependencies, validation steps, and reporting requirements. Plans must be continuously tracked and updated as risk conditions change.

Understand key indicators: KPIs, KRIs, and KCIs. Indicators should be S.M.A.R.T. and aligned with business objectives, risk appetite, and tolerance.

- KPIs measure the performance and effectiveness of controls or processes.
- KRIs provide early warning of changes in risk, based on defined thresholds.
- KCIs measure the effectiveness and reliability of specific controls.

Review Questions

The following questions are designed to test your understanding of this chapter's material. For information on how to obtain additional questions, see this book's Introduction. Answers can be found at the end of this chapter.

1. You are the business process owner for your company's manufacturing department. A risk assessment identifies several data protection risks that could materially affect production operations. Mitigating controls are designed and operated by both the IT security and physical security functions, with performance metrics reported back to the business. Residual risk acceptance authority remains with the business. Which of the following most accurately describes risk and control ownership in this scenario?
   - **A.** You are the risk owner, while IT and physical security are the control owners.
   - **B.** You are the risk owner, and executive management is the control owner due to enterprise governance.
   - **C.** IT and physical security own the risks because they design and operate the controls.
   - **D.** Executive management owns the risks, while you and the security functions share control ownership.

2. As a risk practitioner in a larger organization, you are reviewing risk response options for a risk affecting a mission-critical business process. All feasible and cost-justified controls have been implemented, yet a measurable level of residual risk remains. The risk cannot be transferred to a third party, and discontinuing or materially altering the process would be unacceptable to the business. Executive management must determine the appropriate next action. Which risk response option is MOST appropriate in this situation?
   - **A.** Risk avoidance by redesigning the business process
   - **B.** Risk mitigation through the implementation of additional controls
   - **C.** Risk acceptance with formal approval of residual risk
   - **D.** Risk sharing through contractual or insurance mechanisms

3. Your company has just established a contract with a third-party cloud services provider. Based on the contract language, in the event of a breach, the provider must disclose details of the event as well as allow an audit of the security controls that should have protected any data disclosed during the breach. The contract language, as well as laws and regulations, also stipulate that the third-party provider retains some legal liability for the loss. Which of the following accurately describes the responsibility and accountability borne by both your organization and the third-party provider during this event?
   - **A.** All the responsibility and accountability are placed on the third-party provider.
   - **B.** All of the responsibility and accountability are placed on the organization.
   - **C.** The third-party provider is responsible for protecting the data, but the organization retains accountability.
   - **D.** Both the third-party provider and the organization bear responsibility and accountability for the loss of the data, as well as legal liability.

4. An organization is integrating an emerging technology into its existing IT environment to support a critical business initiative. During planning, management identifies uncertainty related to how the new technology will interact with legacy systems, data flows, and existing controls. Which risk factor represents the MOST significant concern in this scenario?
   - **A.** Adequacy of existing security mechanisms supporting the new technology
   - **B.** Compatibility of data formats between legacy and new systems
   - **C.** Long-term vendor support of the emerging technology
   - **D.** Ability of the new technology to interoperate with existing systems and processes

5. During control testing, a vulnerability scan identifies weaknesses that would normally require patching. However, testing confirms that applying the patches would disrupt several mission-critical line-of-business applications. The vulnerabilities cannot be immediately remediated without significant operational impact. What is the MOST appropriate next step?
   - **A.** Proceed with patching and require business process owners to address any resulting application issues.
   - **B.** Defer patching indefinitely because the affected applications are critical to business operations.
   - **C.** Remove the affected systems from the vulnerability management program until replacement systems are implemented.
   - **D.** Evaluate the residual risk, identify possible compensating controls, and obtain formal management approval for any control exceptions.

6. Your company has implemented several new controls to strengthen security inside its information processing facility. These controls include stronger steel doors with new cipher locks and two additional guard stations to block unauthorized personnel from entering sensitive areas. Which of the following types and functions would those new controls be considered?
   - **A.** Physical, detective
   - **B.** Physical, preventive
   - **C.** Administrative, preventive
   - **D.** Administrative, deterrent

7. Your company is now doing business internationally, and as a result, one of its key business partners overseas requires that it adopt international security standards. You examine several standards and are already using internally developed controls that may be mapped to a new standard. Which of the following control standards is the most appropriate to use when working with international partners?
   - **A.** NIST SP 800-53
   - **B.** CIS controls
   - **C.** ISO/IEC 27001/27002
   - **D.** SOC 2 controls

8. An organization is implementing a new system that processes highly sensitive data. Existing security controls partially address protection requirements, but additional network security devices are being considered to meet heightened data sensitivity obligations under applicable laws and regulations. The organization must ensure that the selected controls both satisfy regulatory requirements and function effectively within the current environment. Which factor is MOST critical when selecting the new security controls?
   - **A.** Cost of the new security devices relative to the existing control environment
   - **B.** Ability of the controls to meet regulatory protection requirements
   - **C.** Degree to which the controls reduce overall enterprise risk
   - **D.** Compatibility of the controls with existing network architecture

9. A new network security device is deployed in parallel with existing infrastructure. Testing confirms that the device meets governance requirements and significantly reduces external attack risk. While most systems function as expected, several legacy and third-party systems experience intermittent connectivity failures when traffic is routed through the device. Which issue should be evaluated first?
   - **A.** Interoperability between the new security device and existing systems
   - **B.** Strength of authentication mechanisms enforced by the device
   - **C.** Adequacy of encryption algorithms supported by the device
   - **D.** Completeness of network protocol hardening configurations

10. Your organization is preparing to implement many new security controls, including upgrading legacy network security devices. While a thorough risk analysis has been conducted, there remains uncertainty about how network devices will interact with the rest of the infrastructure. Management is concerned with interoperability issues and the ability to monitor the installation carefully. Which of the following will help the implementation go smoothly?
    - **A.** Vulnerability assessment prior to installation
    - **B.** Parallel cutover
    - **C.** Risk analysis
    - **D.** Change and configuration management processes

11. You are preparing an annual risk and control effectiveness report for the board of directors using data collected from multiple internal and external sources. To ensure the report supports effective governance and decision-making, certain data characteristics must be evaluated before aggregation and presentation. Which characteristic is LEAST relevant in assessing the suitability of data for inclusion in the report?

A. Completeness of data across all in-scope risk areas

B. Accuracy of data obtained from source systems

C. Reliability of data sources used in the analysis

D. Level of technical detail contained within the data

12. You are creating a schedule for risk reporting for various managers and events. Which of the following strategies should you employ regarding the reporting of risk treatment options and plans?

A. Report on their progress periodically and as needed.

B. Report annually to senior management and quarterly to middle management.

C. Report results only when the risk treatment options have been implemented and are complete.

D. Do not report on risk treatment options; report only on increases or decreases in risk, as required by governance.

13. An organization identifies a significant operational risk related to a core business process supported by IT controls. The risk owner and control owner are assigned to different roles within the organization. Which responsibility is MOST appropriately assigned to the control owner?

A. Determining the organization's risk appetite for the process

B. Accepting residual risk after control implementation

C. Ensuring controls are designed, implemented, and operating effectively

D. Prioritizing risks across the enterprise risk register

14. An organization wants to improve management visibility into changes in overall control effectiveness and risk posture. The reporting approach must support timely awareness of significant changes, present information at an appropriate level for management decision-making, and allow access to more detailed underlying metrics when necessary. Which reporting technique BEST meets these requirements?

A. Heat map summarizing inherent and residual risk levels by risk category

B. Narrative report distributed periodically to executive management

C. Centralized dashboard aggregating key risk and control indicators

D. Scorecard measuring performance against pre-defined targets over time

15. You are developing metrics for senior management on control effectiveness and the organization's risk posture. You need to create a metric that shows managers how the potential risk has been reduced by implementing several new controls over the past six months. You want to aggregate this measurement into one indicator. Which of the following would be the MOST effective indicator to show this reduction?

- **A.** Key performance indicator
- **B.** Key risk indicator
- **C.** Key management indicator
- **D.** Key control indicator

16. An organization is reviewing its control environment after identifying multiple controls that provide overlapping protection with limited additional risk reduction. Management wants to reduce complexity and cost while ensuring continued compliance and effective risk mitigation. Which is the most appropriate first step in a control rationalization initiative?

- **A.** Eliminate redundant controls to reduce operational overhead
- **B.** Map existing controls to risk scenarios, control objectives, and regulatory requirements
- **C.** Replace redundant controls with alternative compensating controls
- **D.** Conduct penetration testing to determine which controls provide the strongest protection

17. A cloud provider suffers a regional outage affecting your organization's systems. You review the shared responsibility model and confirm that the provider was responsible for infrastructure resilience. What is your organization's remaining accountability in this scenario?

- **A.** None; all responsibility transfers to the provider
- **B.** Accountability for data classification, access governance, and continuity planning
- **C.** Only legal liability, but not operational accountability
- **D.** Accountability for the provider's physical infrastructure controls

18. During risk reporting, you notice that one KRI has crossed its defined threshold for three consecutive reporting periods. Which action should occur next?

- **A.** Escalate the condition and initiate pre-defined response procedures
- **B.** Automatically accept the increased risk as unavoidable
- **C.** Replace the KRI with a KPI to improve performance tracking
- **D.** Perform a full enterprise risk assessment before acting

19. A business unit requests an exception to a mandatory encryption control due to operational constraints affecting a time-critical system. The exception is approved as part of the organization's formal exception management process. Senior management requires assurance that the exception does not weaken governance or obscure accountability over time. Which combination of elements is MOST critical to include in the exception documentation to ensure effective risk oversight?

- **A.** Business justification for the exception and acknowledgment of increased risk by the requestor
- **B.** Defined expiration date, documented residual risk acceptance authority, and required interim mitigation measures

C. Statement that the exception applies only to the requesting business unit and cannot be reused elsewhere

D. Description of the control objective and confirmation that the exception aligns with enterprise risk appetite

20. Your organization is deploying a new AI-enabled analytics system. During the emerging risk review, which factor represents the most common risk associated with integrating new technologies into an existing environment?

A. Vendor licensing model

B. Interoperability and compatibility with existing controls and systems

C. Number of developers trained on the platform

D. The provider's marketing claims regarding system capabilities

21. An organization has identified a high inherent risk related to a critical business application. Management has approved a risk treatment plan that includes implementing additional controls and accepting a defined level of residual risk. Which activity is MOST important to include in the risk treatment plan to ensure its ongoing effectiveness?

A. Documenting the organization's risk appetite statement

B. Assigning risk ownership and defining monitoring metrics

C. Updating the enterprise risk register annually

D. Performing a one-time cost-benefit analysis of the selected controls

22. An enterprise uses multiple internal and third-party data sources to assess key risk indicators (KRIs) supporting risk response decisions. During a review, management discovers that risk reports are inconsistent across business units despite using the same underlying datasets. Which weakness is MOST likely the primary cause of this issue?

A. Lack of formally defined risk appetite thresholds

B. Insufficient frequency of data collection from source systems

C. Failure to align KRIs with business objectives

D. Inadequate validation of data transformation and aggregation logic

23. An organization implements new security controls as part of a risk response initiative. Management wants to evaluate whether the controls are operating as intended and delivering expected outcomes. Which key performance indicator (KPI) is MOST appropriate for this purpose?

A. Percentage of implemented controls meeting defined performance targets

B. Number of risk events reported during the last reporting period

C. Estimated financial impact of potential future risk scenarios

D. Number of audit findings related to policy non-compliance

24. An organization has implemented compensating controls to reduce residual risk associated with a critical business process. Management wants early warning signals that these controls may be weakening before a risk event occurs. Which metric BEST serves this purpose as a key control indicator (KCI)?
    - **A.** Percentage of business objectives achieved during the quarter
    - **B.** Total financial losses resulting from realized risk events
    - **C.** Overall customer satisfaction score for the process
    - **D.** Number of control exceptions identified during routine monitoring

Answers to Review Questions

1. A. As the business process owner, you are accountable for the risks affecting the manufacturing process and for accepting residual risk. IT security and physical security functions design, implement, and operate controls, making them the control owners. Ownership is not determined by who implements safeguards, but by accountability for risk outcomes.

2. C. When a risk affects a critical business process, cannot be transferred, and all reasonable mitigation measures have already been applied, the remaining exposure is residual risk. In this case, the appropriate response is formal risk acceptance by authorized management, supported by documentation and ongoing monitoring.

3. D. Both the third-party provider and the organization bear some level of responsibility, accountability, and legal liability for the data loss, as the contract language and laws address this situation. However, it is likely that the organization will ultimately bear the most responsibility and accountability for the data loss to its customers.

4. D. For emerging technologies, interoperability is a primary risk because incompatibilities with legacy systems, interfaces, or control environments can disrupt operations, weaken controls, and introduce unanticipated risks. This uncertainty is a defining characteristic of emerging risk.

5. D. You should evaluate the risk of patching systems and suffering the impact to the line-of-business applications versus the security risk if the systems are not patched, and then allow management to decide based on that risk. Any exceptions should be formally approved and documented. You cannot simply ignore the risk. Taking any action without carefully considering it would be detrimental to the organization, whether it is choosing to break line-of-business applications or ignoring the security ramifications of the problem.

6. B. The controls are physical because they use physical barriers to secure the facility's interior. The controls are also preventive because they prevent unauthorized entry into sensitive, restricted areas. The controls are not administrative because they do not involve written policy or procedures established by management.

7. C. ISO/IEC 27001 and 27002 are international standards developed to ensure interoperability and translation of security controls across international boundaries. The other control sets, while they may be used in areas outside the United States, are unique to the United States.

8. B. When laws or regulations impose specific protection obligations for sensitive systems or data, those requirements establish a non-negotiable baseline for control design and selection. Controls that do not meet regulatory protection requirements are unacceptable, regardless of their cost efficiency, interoperability, or overall risk reduction benefits. Therefore, the primary consideration in selecting new security controls is their ability to satisfy applicable regulatory obligations.

9. A. Selective connectivity failures affecting only certain systems—particularly legacy or third-party platforms—are most commonly caused by interoperability issues. These can include incompatibilities in protocol handling, interface standards, traffic inspection methods, or unsupported features. Interoperability directly affects a control's ability to function effectively within an existing environment, even when governance, compliance, and risk-reduction objectives are met.

10. D. Proper change and configuration management processes will assist in an orderly implementation and change to the new network security devices. Management can carefully consider, test, and monitor infrastructure changes, as well as back out changes according to defined processes if something goes wrong.

11. D. Completeness, accuracy, and reliability directly affect the integrity and credibility of risk and control reporting. The level of technical detail, however, is a presentation concern, not a data suitability concern. Technical complexity can be managed during summarization and reporting without compromising data quality.

12. A. You should report the progress of risk treatment options both periodically, as determined by organizational policy, and as needed, depending on the criticality or importance of the information and the desires of management.

13. C. The control owner is responsible for the design, implementation, and ongoing operation of controls that mitigate identified risks. This includes maintaining control effectiveness and addressing control deficiencies.

14. C. A centralized dashboard supports effective risk and control reporting by aggregating key risk indicators (KRIs) and key control indicators (KCIs) into a single, continuously updated view for management. Dashboards are designed to highlight material changes in risk posture and control effectiveness, enabling timely decision-making without overwhelming executives with unnecessary detail. Unlike heat maps or scorecards, dashboards allow users to drill down into underlying metrics and source data as needed, while maintaining a high-level, management-focused perspective. This combination of timeliness, aggregation, and selective detail makes a centralized dashboard the most appropriate reporting technique in this scenario.

15. B. In this scenario, a key risk indicator (KRI) could be developed to show the risk reduction. A key performance indicator (KPI) would measure the performance of a control or other aspect of organizational security, not necessarily the risk reduction. A key control indicator (KCI) would show the effectiveness of a particular control, not the risk reduction. Key management indicators are published by executive management and cover various metrics they are concerned with regarding the overall performance of organizational objectives. They do not indicate a risk reduction.

16. B. Control rationalization begins with a structured analysis that maps existing controls to defined risk scenarios, control objectives, and applicable regulatory or framework requirements. This analysis establishes whether controls are necessary, duplicative, or misaligned with current risk exposure. Only after understanding how each control contributes to risk mitigation and compliance should management decide whether controls can be removed, consolidated, or replaced without increasing residual risk.

17. B. In shared responsibility models, the customer organization retains accountability for data governance, identity and access management, and continuity planning. Infrastructure resilience is the cloud provider's responsibility, but customer-level governance obligations remain. In this scenario, the customer organization should establish an alternate processing plan, such as a separate availability zone, a different cloud provider, or on-premise processing.

18. A. KRIs with breached thresholds require pre-defined escalation, investigation, and treatment steps. They act as early warnings and must trigger a response, not automatic acceptance or replacement.

19. C. Effective exception management requires that exceptions are temporary, controlled, and accountable. A defined expiration date ensures the exception is revisited, documented residual risk acceptance establishes clear ownership and authorization, and interim mitigation measures reduce exposure while the control is not fully enforced. Together, these elements preserve governance, enable monitoring, and prevent exceptions from becoming de facto permanent control failures. Exceptions must be time-bound and reassessed periodically, because this year's acceptable risk and accompanying exception may not be acceptable in future years.

20. B. Interoperability and compatibility risks are the most significant concerns with emerging technologies—affecting integration, data flows, security controls, performance, and reliability.

21. B. A risk treatment plan must clearly assign ownership and define metrics to monitor control performance and residual risk over time. This ensures accountability, supports ongoing oversight, and enables timely adjustments if the risk posture changes.

22. D. Inconsistent risk reporting across business units using the same datasets most commonly results from weaknesses in data transformation, aggregation, or normalization processes. Without proper validation controls, data may be interpreted or combined differently, leading to inconsistent analysis and unreliable risk response decisions.

23. A. KPIs measure the performance and effectiveness of processes and controls. Tracking the percentage of controls meeting defined performance targets directly assesses whether risk response actions are operating as intended and achieving expected results.

24. D. KCIs are designed to provide early warning of control effectiveness degradation. Tracking control exceptions directly reflects whether controls are failing or being bypassed, allowing management to intervene before a risk event materializes.



Chapter 4

Information Technology

This chapter covers CRISC Domain 4, "Information Technology and Security," specifically the Information Technology Principles subdomain. The job practice for this domain is described as follows:

In this domain, we interrogate the alignment of business practices with Risk Management and Information Security frameworks and standards, as well as the development of a risk-aware culture and implementation of security awareness training.

The preceding statement from ISACA describes the whole of Domain 4. This chapter covers the information technology portion of this domain, while Chapter 5 covers the information security portion.

This chapter covers these job practice elements:

✔ A—INFORMATION TECHNOLOGY PRINCIPLES

1. *Enterprise Architecture*
2. *IT Operations Management (e.g., change management, IT assets, problems, incidents)*
3. *Project Management*
4. *Disaster Recovery Management (DRM)*
5. *Data Lifecycle Management*
6. *System Development Lifecycle (SDLC)*
7. *Emerging Technologies*

There are no CRISC Task Statements associated with this domain.

The topics in this chapter and Chapter 5 represent 20% of the CRISC examination.

Even if an organization is in the business of producing goods to bring to market, it still relies on its information and its technology as enablers to produce those goods and deliver them to the market and

consumers. Without the ability to generate, process, and otherwise use information, modern businesses would not be able to produce goods or services and compete in market spaces. Therefore, any negative events that affect the organization's ability to process information directly affect the organization's ability to survive. Information technology directly supports business goals, objectives, and strategies. All elements of a business endeavor depend on IT, including the organization's information (or data), its people, its line-of-business applications and systems, and its overall infrastructure, including all of its equipment and processes.

Information technology poses risks to the business enterprise in two different ways. First, information technology is used to protect the enterprise from risk; in other words, it safeguards the information the business generates and relies on. That's the purpose of having high-capacity storage, faster networking equipment, redundant systems, and security devices sprinkled throughout the infrastructure. The second way IT affects risk is by enabling the organization to produce the information needed to achieve its business goals and objectives. Without IT, there would be no information processed, no systems designed, and no advanced technologies developed for a business to take to market and compete with. Therefore, IT serves to both protect information and generate it to advance business goals.

The information the business generates and uses has several key characteristics. First, the information should be relevant to the processes that the business supports. It should also be timely; stale information can prevent a business from fulfilling its functions in both short-term and long-term timeframes. It should also be accurate and complete.

Enterprise Architecture (EA)

An organization's enterprise architecture affects its risk in several ways. Aspects of enterprise architecture risk include interoperability, supportability, security, maintenance, and how the components of the infrastructure fit within the systems development lifecycle. The business views IT as an investment of capital funds, much like it does facilities and other equipment, as a means of supporting the business mission. Information systems pose risks to the business due to interoperability, supportability, security, and other factors. It costs the organization money to maintain and support all IT assets, including parts, software licenses, training for administrators and users, and upgrades. There are also the intangible aspects of IT, such as business value and liability. IT systems affect the bottom line of the organization, so a lot of thought is put into managing risk for them. Additionally, you should remember that information technology risk is only one component of the enterprise risk picture. The next few sections discuss various aspects of information systems architecture and explain how they contribute to the organization's overall enterprise risk.

EA Concepts and Structures

Enterprise architecture (EA) provides a structured approach for describing how an organization's business processes, information flows, applications, and technology infrastructure work together to support strategic objectives. EA establishes common

principles, standardized reference models known as *reference architectures*, and defined architectural layers that guide system design, technology selection, and lifecycle governance. By documenting relationships among business functions, data assets, applications, and platforms, EA enables consistency, interoperability, and controlled evolution of the environment. The subsections that follow examine the fundamental structures and components of EA and how they shape risk, cost, and operational performance.

Business–technology alignment is the degree to which an organization's information technology capabilities support, enable, and extend its business strategy, objectives, and operating model. Effective alignment ensures that technology investments, architectures, processes, and services are directly tied to business needs and risk priorities, rather than operating as independent or siloed functions.

Alignment begins with recognizing that technology must be governed and managed as an integrated component of enterprise value delivery. To achieve this, organizations define clear relationships among enterprise strategy, business processes, information flows, and the underlying technology architecture. Misalignment results in increased operational risk, uncontrolled technology proliferation, technology debt, inconsistent controls, and gaps in service delivery.

Several principles support effective business–technology alignment:

- **Strategic alignment:** Strategic alignment requires that IT initiatives be driven by documented business goals and measurable outcomes. Technology roadmaps, enterprise architecture models, and IT portfolios must be evaluated against strategic objectives to ensure that projects and services contribute to business performance, resilience, regulatory compliance, and acceptable risk levels.
- **Process alignment:** Business processes rely on technology to execute workflows, manage information, and deliver services. Technology design should reflect process requirements for availability, integrity, confidentiality, scalability, and continuity. Misalignment can cause inconsistent workflows, duplicated systems, and inefficiencies that directly increase enterprise risk.
- **Information alignment:** Business information requirements determine how data is generated, classified, stored, transmitted, and protected. Alignment requires clear mapping between business information needs and the supporting applications, databases, and infrastructure. When information flows do not align with business requirements, risks arise in data quality, accessibility, and regulatory compliance.
- **Architectural alignment:** Enterprise architecture provides a structure for how business, data, application, and technology layers interconnect. Alignment requires that architecture decisions reflect business constraints, service-level needs, budget requirements, and operational capabilities. Poor architectural alignment introduces interoperability issues, support challenges, and lifecycle management risks.
- **Risk alignment:** Business–technology alignment includes evaluating how technology decisions affect enterprise risk posture. Technology architectures must reflect risk appetite, risk tolerance, regulatory obligations, and control requirements. When technology evolves more rapidly than business risk management processes, the organization may experience control gaps, inconsistent protections, or unmanaged exposure.

- **Governance alignment:** Alignment depends on governance structures that ensure coordinated decision-making. This includes clear definitions of roles and responsibilities, oversight mechanisms for technology investments, and accountability for lifecycle management. Governance structures must enable traceability from business decisions to technology outcomes, including financial, operational, and security impacts.
- **Lifecycle alignment:** Technology must be governed throughout its lifecycle—from planning and acquisition to operation and retirement—based on documented business needs. Alignment ensures that lifecycle activities such as change management, configuration management, DR/BC planning, and SDLC processes consistently support business performance and risk objectives.

Business-technology alignment is depicted and reported in several ways, including IT balanced scorecards and metrics (including KPIs and KRIs), and capability (maturity) models.

Effective business–technology alignment results in predictable service delivery, optimized technology investments, reduced operational and security risks, and improved ability to adapt to changes in the business environment. Poor alignment increases complexity, cost, technology debt, and exposure to operational disruptions and regulatory non-compliance.

Enterprise Architecture Frameworks

Another facet to consider is enterprise architecture. This differs from network architecture presented earlier. EA focuses on the relationship between business and technology. In most cases, this includes the relationship between security and technology. It is very important because it provides a view into the governance of an organization, which can provide insight into some core problems that cause risks within the organization.

Two frameworks for enterprise architecture are discussed in this book: the Open Group Architecture Framework (TOGAF) and the Zachman Framework.



TOGAF and the Zachman Framework are enterprise architecture models, not enterprise *security* architecture models.

The Open Group Architecture Framework

The *Open Group Architecture Framework (TOGAF)* is a lifecycle enterprise architecture framework used for designing, planning, implementing, and governing an enterprise technology architecture. TOGAF could be considered a high-level approach for designing enterprise infrastructure.

The phases used in TOGAF are as follows:

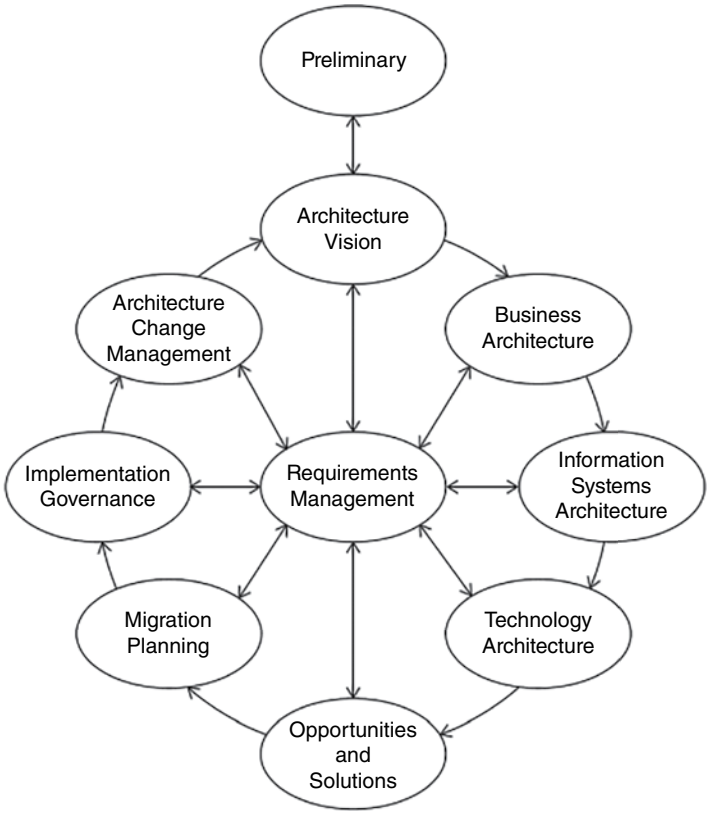
- Preliminary
- Architecture vision
- Business architecture

- Information systems architecture
- Technology architecture
- Opportunities and solutions
- Migration planning
- Implementation governance
- Architecture change management
- Requirements management

TOGAF is a business-driven, lifecycle management framework for enterprise architecture overall, and it certainly can be used for information security architecture as well. Figure 4.1 depicts TOGAF visually.

You can find information on TOGAF at www.opengroup.org/togaf.

FIGURE 4.1 TOGAF components.



Source: Adapted from https://www.opengroup.org/togaf.

The Zachman Framework

The *Zachman Enterprise Architecture Framework*, established in the late 1980s, remains a dominant EA standard today. Zachman likens IT enterprise architecture to the construction and maintenance of an office building: at an abstract level, the office building performs functions such as containing office space. As you look into increasing levels of detail in the building, you encounter various trades (steel, concrete, drywall, electrical, plumbing, telephone, fire control, elevators, and so on), each of which has its own specifications, standards, regulations, construction and maintenance methods, and so on. Zachman can be used to develop an enterprise security architecture within the enterprise architecture.

In the Zachman Framework, IT systems and environments are described at a high, functional level and then in increasing detail, encompassing systems, databases, applications, networks, and so on. Table 4.1 illustrates the Zachman Framework.

While the Zachman Framework allows an organization to examine cross-sections of an IT environment that supports business processes, the model does not capture the relationships among IT systems. Data flow diagrams are used instead to depict information flows.

Information about the Zachman Framework is available at www.zachman.com/about-the-zachman-framework.

EA Risk Considerations

Aside from cybersecurity risks, which are discussed in Chapter 5, IT organizations need to be aware of certain operational risks associated with enterprise architecture. These risks include:

- **Lack of consistency:** Organizations lacking architectural discipline are likely to grow their environments organically, which will result in a lack of consistency. With discipline, organizations can develop *reference architectures*, which are detailed designs of portions of an environment. With reference architectures, organizations will implement technology according to the reference architecture, resulting in a consistent approach to various business needs.
- **Lack of documentation:** Organizations can manage their infrastructure only to the extent that they understand what it is, how it works, and how it is managed. Documentation, in the form of architecture, maintenance, and operations procedures, ensures that the organization can better manage the environment and experience less unplanned downtime.
- **Technical debt:** Organizations need to track all their hardware, software, and information assets, and in so doing, track the status of vendor support for each component. Even so, organizations may find themselves with outdated, undocumented, and unsupported infrastructure, known as *technical debt*.

TABLE 4.1 Zachman Framework showing IT systems in increasing levels of detail.

	Data	Functional (Application)	Network (Technology)	People (Organization)	Time	Strategy
Scope	List of datasets important in the business	List of business processes	List of business locations	List of organizations	List of events	List of business goals and strategy
Enterprise Model	Conceptual data/object model	Business process model	Business logistics	Workflow	Master schedule	Business plan
Systems Model	Logical data model	System architecture	Detailed system architecture	Human interface architecture	Processing structure	Business rule model
Technology Model	Physical data/ class model	Technology design	Technology architecture	Presentation architecture	Control structure	Rule design
Detailed Representation	Data definition	Program	Network architecture	Security architecture	Time definition	Rule speculation
Function Enterprise	Usable data	Working function	Usable network	Functioning organization	Implemented schedule	Working strategy

Source: Adapted from The Zachman Framework.

IT Operations Management

Managing enterprise infrastructure is among the most labor-intensive tasks in an organization, particularly in larger organizations. The enterprise infrastructure in an organization covers a broad scope of areas. Infrastructure, of course, covers workstations, printers, and other end-user devices. It also covers the major pieces needed to conduct business, such as servers, cabling, switches, routers, and a variety of other network and security devices. Most organizations also include wireless networks in their supporting equipment. Additionally, the infiltration of mobile devices into businesses makes this a new and sometimes difficult area to bring under the organization's infrastructure umbrella. Challenges with enterprise infrastructure include not only maintenance, upgrades, updates, and the implementation of new technologies, but also more traditional management challenges, such as budgeting, project management, and staffing.

Key areas in IT operations management include server and infrastructure management, end-user support, help desk, problem escalation management, and, of course, cybersecurity. IT managers must maintain and meet any internal agreements that cover guaranteed service delivery in support of the different divisions within the organization, as well as service-level agreements (SLAs) with other organizations. While information security personnel tend to focus on IT risk, we find that most risk factors, threats, and vulnerabilities affecting IT also affect operations across the business. In the next few paragraphs, however, I focus on those that directly affect managing the IT operations in an organization.

Management of IT operations incurs risk factors such as the size and complexity of the organization and its IT infrastructure, the criticality and priority of the different systems that make up the infrastructure, and the internal management processes of the organization. Resourcing issues also affect the management component of IT operations; staffing the right people who are trained on the various technologies the organization needs, as well as creating a budget that helps maintain the current level of support and support future growth for IT, are examples of resource-related risk factors. Additionally, compliance with legal and governance requirements in the face of increased regulation is also a risk factor that IT managers must deal with.

Organizational structure also affects IT operations because the IT infrastructure may be managed on either a centralized or a decentralized basis—sometimes by one centralized IT shop, or in the case of larger organizations, by multiple areas within other divisions that are delegated the task of managing their own piece of the infrastructure. IT might also be managed on a functional basis; the accounting and human resources divisions may have the responsibility for managing the systems that support their mission. Additionally, the IT infrastructure may be managed on a geographical basis, in the case of physically separated locations in large organizations. All these structural considerations contribute to risk in managing IT operations.

Also directly affecting IT operations risk are two key processes: the change and configuration management processes. I've already discussed risk factors that affect the SDLC model in an organization, as well as how new and emerging technologies can affect the existing infrastructure. These same risk factors also directly affect the daily management of

IT operations simply because they introduce change into the network. Some of this change can be planned and carefully introduced. Still, often there is change—from the large-scale network side, all the way down to individual configuration items on hosts—that may have unforeseen consequences and affect the network in different ways. How the organization addresses change and configuration management is a risk factor that affects not only the SDLC but also day-to-day operations.

Threats that affect the IT operations are like those that affect other areas of the business and could be external or internal threats from a variety of sources. Some of these threats have a ripple effect in that they may first affect other areas of business and, in turn, affect IT. For example, the threat of a poor economy may affect the profitability and sustainability of the business, resulting in cuts, which often include IT personnel or equipment. Obviously, there are also direct threats to the daily management of the IT infrastructure. These include external threats such as hackers, of course, but also more commonly come as increases in cost to maintain operations, issues with external service providers, and disaster-related events (weather, fire, and so on). Internal threats could include those carried out by malicious insiders or even careless workers, such as theft, sabotage, accidental equipment breakage, and so on. Internal threats could also come from management processes and include budget cuts, loss of trained personnel, shifts in organizational priorities, and transitioning from one market segment to another.

Technical vulnerabilities are beyond the scope and purpose of this book; our focus here is on those associated with the organization's IT management. Vulnerabilities affecting IT management include faulty processes and procedures, insufficient resources, and a shortage of trained technical personnel. An end-user population that has not been trained or held accountable for the proper care and treatment of the infrastructure may also pose a vulnerability. Other vulnerabilities might include a lack of infrastructure monitoring, insufficient control over sensitive information in IT systems, and failure to maintain a stable SDLC across the entire infrastructure.

Core Operations Functions

Core operations functions encompass the day-to-day activities required to sustain the availability, performance, and reliability of an organization's technology environment. These functions implement the operational controls, processes, and procedures defined by IT governance and enterprise architecture, to ensure that systems and services remain stable, secure, and aligned with business requirements. Operational functions also serve as key control points for identifying, managing, and mitigating risk within IT services.

Two well-known international standards define a suite of *IT service management* activities, which are:

- ***IT Infrastructure Library (ITIL)*:** Initially developed in the UK in the 1980s to standardize primary IT business processes.
- ***ISO/IEC 20000*:** An international standard that borrowed heavily from ITIL.

Core IT functions related to ITIL and ISO/IEC 20000 are described next.

Infrastructure and Platform Operations

Infrastructure operations include the deployment, configuration, monitoring, and maintenance of servers, storage systems, networks, connected devices, and cloud resources. These activities ensure that compute, network, and storage platforms maintain required performance levels and meet defined service-level agreements (SLAs). Core tasks include asset management, lifecycle management, capacity planning, patching, routine maintenance, and response to infrastructure-related incidents or disruptions. Modern operations extend these responsibilities to virtualized environments, cloud platforms, and container orchestration systems, requiring standardized configurations and automated management practices.

End-user Services and Device Management

End-user services support the desktops, laptops, mobile devices, peripherals, and productivity tools used throughout the organization. Core functions include provisioning, configuration, asset tracking, software distribution, patching, and security control enforcement. Centralized device management platforms allow consistent application of security baselines, configuration standards, event detection, and access controls. Effective end-user support reduces downtime, limits configuration drift, and ensures compliance with policies related to data handling, authentication, and endpoint protection.

IT Service Desk and Incident Management

The *IT service desk* provides a centralized point of contact for users to report issues, request services, and obtain support. Incident management processes ensure that disruptions are logged, categorized, prioritized, and resolved in a controlled manner. Core responsibilities include intake, triage, initial analysis and diagnosis, escalation, communication to stakeholders, and documentation of corrective actions. Well-defined *incident management* processes reduce service downtime, enable trend analysis, and support compliance with service-level and risk management requirements.

Problem Management

Problem management identifies the root causes of recurring or significant incidents and implements permanent fixes. Activities include incident trend analysis, *root cause analysis (RCA)*, evaluation of systemic failures, and implementation of corrective or preventative actions. By addressing underlying causes rather than symptoms, problem management reduces operational risk, minimizes repetitive disruptions, and supports long-term service stability.

Change Management

Change management processes control modifications to infrastructure, applications, configurations, and services. The process ensures that proposed changes are assessed for business impact, risk, dependencies, resource requirements, and rollback procedures. Approved changes are scheduled, implemented, verified, and documented. A disciplined change management function reduces the likelihood of unscheduled service outages, configuration errors, and unintended security impacts introduced through poorly controlled modifications.

Configuration Management

Configuration management maintains authoritative records of system components, their attributes, and their relationships. *Configuration management databases (CMDBs)* and related tools capture information about hardware, software, dependencies, and configuration baselines. This information supports effective troubleshooting, impact assessment, vulnerability management, and lifecycle planning. Configuration management also helps enforce standard builds and reduce drift across on-premises and cloud environments.

Monitoring, Logging, and Observability

Continuous monitoring and observability functions track system health, performance, availability, and security telemetry. Core activities include infrastructure monitoring, application performance monitoring, log collection, alerting, and threshold- or behavioral-anomaly detection. Observability platforms ingest metrics, logs, and traces from distributed or cloud-native workloads, enabling early detection of degradations, capacity constraints, and operational risks. Effective monitoring supports incident response, performance tuning, and compliance with resilience and availability objectives. Most organizations employ a security information and event management system (SIEM), a *threat intelligence platform (TIP)* to keep the SIEM informed of emerging threats, and a *security orchestration and response (SOAR)* system to automate responses to common security events.

Backup, Recovery, and Data Protection

Operational teams ensure implementation of backup schedules, verification of backup integrity, and maintenance of restore capabilities. Activities include management of backup repositories, offsite or immutable storage mechanisms, and periodic testing of restore procedures. These functions directly support business continuity and disaster recovery objectives by enabling recovery of systems and data following failure, corruption, cyber incidents, or accidental deletions. Operational control over data protection mitigates risks associated with data loss, ransomware, and service disruption. In response to *ransomware* threats (where criminals encrypt data and demand a ransom in exchange for a decryption key), many organizations employ *immutable backups* (where the backup medium cannot later be altered or destroyed) to enhance their ability to recover from such attacks.

Identity and Access Operations

Identity operations include account provisioning, de-provisioning, role management, credential lifecycle management, and enforcement of authentication and authorization policies. These activities ensure that access rights reflect current job responsibilities, that orphaned accounts are eliminated, and that authentication mechanisms, such as MFA, privileged access management, and identity federation, remain operational and aligned with organizational policy. Misalignment in identity operations introduces significant security and compliance risks.

Vendor and Service Provider Coordination

Many organizations rely on external service providers for network connectivity, cloud services, SaaS applications, managed security services, and specialized infrastructure. Core operational responsibilities include coordinating incidents, changes, and service requests with providers; validating provider performance against contractual obligations; and monitoring compliance with security and operational requirements. Poor coordination or oversight can introduce dependency risks, visibility gaps, or uncontrolled service disruptions.

Lifecycle and Asset Management

Asset management tracks the acquisition, deployment, maintenance, and retirement of hardware and software. Effective asset management reduces technical debt, supports secure configurations, and ensures that unsupported or obsolete components do not introduce operational or security risks. Core activities include:

- **Inventory management:** Asset management is critical for several IT and cybersecurity functions, including vulnerability management, incident management, license management, and more. Key activities to ensure accurate and complete inventory include *asset discovery*, which uses various methods to ensure that IT is aware of all components residing on networks, in systems, and in data stores. Asset discovery tools can also be used to detect *shadow IT*, the unauthorized use of hardware and software to support business processes.
- **Licensing compliance:** Organizations must keep track of the use of software licenses, to ensure compliance with contractual terms. Often, organizations are free to implement more copies of software than license agreements permit, leading to out-of-compliance situations and possible fines or legal action.
- **Refresh and end-of-support planning:** Organizations need to carefully track the timespan for vendor support of every hardware and software component and plan for the refresh of components prior to their end of life.

Security Management

Security management is included as a part of IT service management and is included in ITIL and ISO/IEC 20000. However, in the CRISC job practice, security is in the second part of the Information Technology and Security domain, and it is discussed fully in Chapter 5.

Operations Governance and Reporting

Operational governance ensures that processes follow documented procedures, controls are consistently applied, and results are measured against defined performance indicators. Reporting activities track service availability, incident trends, change success rates, configuration compliance, and operational risks. These metrics support informed decision-making, capacity planning, and accountability for service performance.

Operational Risk

Operational risk refers to the potential for loss, disruption, or reduced performance resulting from failures in processes, technology, people, or external dependencies that support IT operations. These risks arise from the daily management of infrastructure, applications, services, and supporting functions. Effective identification and control of operational risk is essential to maintaining service continuity, meeting business objectives, and preserving the organization's risk posture. Note that operational risk does not necessarily come directly from IT infrastructure or assets; it is the risk that is inherent to business processes and their interactions with supporting systems, data, and applications. For instance, poor security practices in a business process itself generate operational risk.

Process-related Risks

Process risks occur when operational activities are inconsistent, incomplete, or poorly controlled. Weak or undocumented procedures, inadequate change management, insufficient configuration control, and ineffective incident and problem management contribute to operational instability. A lack of standardized workflows or reliance on ad hoc practices increases the likelihood of outages, misconfigurations, and delayed recovery from service disruptions.

Technology-related Risks

Technology risks arise from failures or weaknesses in systems, platforms, networks, and supporting components. These risks include hardware failures, software defects, outdated operating systems, unsupported products, and technology debt. Rapidly evolving environments—such as cloud platforms, containerized workloads, and distributed architectures—can introduce interoperability issues, configuration drift, and complexity that exceed operational capacity. Poor monitoring or inadequate observability reduces the ability to detect failures before they affect service delivery.

People-related Risks

People risks result from human error, skill gaps, insufficient training, or inadequate staffing. Incorrect changes, improper maintenance actions, failure to follow procedures, and reliance on a small number of key personnel can all impair operational stability. Personnel turnover, lack of succession planning, and limited technical expertise increase dependency risks and can hinder timely response to incidents. Privileged access misuse, whether accidental or intentional, also represents a significant operational exposure. People-related risk is commonly referred to as *insider threat*.

Security-related Operational Risks

Security-related risks include operational failures that expose systems to unauthorized access, data loss, or service disruption. Misconfigured identity systems, missing patches, weak endpoint controls, and inconsistent security baselines can create vulnerabilities exploitable by attackers. Inadequate coordination between operations and security functions may result in incomplete remediation, delayed response to threats, or controls that are not aligned with operational capabilities.

Vendor and Third-party Risks

Dependence on cloud providers, managed service providers, SaaS platforms, and telecommunications vendors introduces risks related to service outages, contractual limitations, and reduced visibility into operational issues. External changes, such as provider maintenance, configuration updates, or policy shifts, can impact internal service availability. Poor oversight, inadequate SLAs, or failure to validate provider performance can lead to unmanaged operational exposures. *Third-party risk management*, discussed in Chapter 3, comprises business processes to identify and manage risks related to vendors and third parties.

Environmental and Physical Risks

Environmental risks include facility outages, power failures, cooling issues, and physical incidents that impair data centers or network locations. Even in cloud-centric environments, physical dependencies, such as local networking, connectivity, and user endpoints, can create operational disruptions. This chapter and Chapter 5 cover disaster recovery planning and business continuity planning, respectively.

Governance and Compliance Risks

Weak governance contributes to operational risk when roles, responsibilities, and decision-making authority are unclear. Insufficient oversight, ineffective control monitoring, or a lack of performance metrics can lead to unmanaged issues and compliance failures. Misalignment between operational practices and policy requirements increases the likelihood of audit findings, regulatory impact, or control deficiencies.

Infrastructure Risk

The implementation and management of various IT infrastructure layers pose risks that enterprise architects must understand. IT infrastructure is growing in its complexity, and architects need to understand how it all works—in general, how it has been implemented and how it is operated and maintained.

Networks

Networks are another aspect of enterprise infrastructure that is essential for facilitating communication between systems. They are most effective when they are implemented in a way that not only incorporates the business logic of the organization but also follows the data flow, considering how the software is most often implemented and integrated with other systems, such as databases.

Let's first take a step back to understand what a network is. At its most basic level, a network is a mechanism that allows systems to communicate with each other. The components are often devices such as *switches* (for local network communication) and *routers* (for communications between networks). A variety of security devices may be connected to a network, including firewalls and intrusion prevention systems, which are discussed in Chapter 5.

In the end, networks and networking can be highly challenging and detailed. What I have covered here only scratches the surface of what is entailed by networks and networking. It is worth your time to explore this topic in depth to help create a better network and ultimately add to a risk assessment.



The more a risk manager understands the implementation details of information security and information technology, the better they will be able to identify risks.

Operating Systems

Although I discussed platforms in a previous section, it's also worth mentioning operating systems as their own separate risk element in the enterprise infrastructure. The terms *platform* and *operating system* are sometimes used interchangeably, but, truthfully, a platform is more a hardware architecture than an operating system categorization. A platform could be an Intel PC or a tablet chipset, for example, which are designed and architected differently and run totally different operating systems.

Different operating systems, on the other hand, could run on the same platform but still introduce risk into the organization for the same reasons discussed previously with applications. For example, there are interoperability and supportability risks, and all the other issues that go hand in hand with the normal operating system lifecycle, such as patch and vulnerability management. Licensing, standardization, level of user control, and configuration are also issues that introduce risk into the organizational computing environment.



Although information security professionals tend to focus on the technical aspects of IT risks, for the exam, keep in mind that all IT risks also contribute to the enterprise's overall business risks. Make sure you look at the larger picture beyond the IT realm.

Gateways

Several other technologies that can be used to protect organizations can be factored into a risk assessment. For example, secure email *gateways* and secure web gateways are a must-have to protect organizations from web-based threats. These technologies can help filter malware, block malicious links or websites, help prevent phishing, or perform a host of other tasks. Technically, they do not fit into the category of other technologies, so they deserve special mention. A good risk assessor will learn about these technologies and how they can help reduce organizational risks. That knowledge can be applied to the risk assessment process.

Software

Software introduces risk simply because it is so critical to business operations today. Businesses need not only basic word processing and spreadsheet software but also complex databases, line-of-business applications, specialized software, security software, and other

types of applications. Software must be managed within its own lifecycle as well; it is constantly being patched, upgraded, superseded, and replaced by better, faster software with more features that usually costs more. Risks that are inherent to managing applications within an organization include supportability, backward compatibility, data format compatibility, licensing, and proper use.

Adding to this complexity are the decisions an organization makes regarding the selection of proprietary software, open-source software, general-purpose commercial software, or highly specialized software. These categories differ in cost, supportability, licensing, and feature sets. Interoperability also contributes to application risk, as it does with other infrastructure components. Applications that do not use common data formats or produce usable output for the organization create the risk of expense or additional work that goes into transforming data between incompatible applications. Applications also introduce risk into the business environment, with the level of security mechanisms built into them and how effective those security mechanisms are in protecting data residing in the application.

It's worth mentioning here that web-based applications, in addition to presenting the same risks as normal client-server apps, also have their own unique risks. Web-based applications introduce risks since they often directly connect to unprotected networks such as the Internet. Other risks include those that come from the wide variety of web programming languages and standards available for developers to use.

Databases

Databases impose some of the same risks that applications and other software do. Additionally, databases incur risks associated with *data aggregation*, compatibility, privacy, and security. Aggregation and inference are risks associated with database systems.

More often than not, databases are actively managed by *database management systems (DBMSs)*, which are a collection of software programs used to control access, modifications, and logging of events concerning one or more databases. Examples of DBMSs include Microsoft SQL Server, Oracle, MySQL, and PostgreSQL. Unauthorized access and data loss are significant risks that databases pose to the enterprise environment.

Platforms

Platforms are components of IT infrastructure that contribute to information security and business risk for several reasons. First, it is costly to field and simultaneously maintain different operating systems and environments across different platforms. Platforms also introduce risk into the environment in the form of interoperability, security, and supportability. A diverse platform environment (with mixed platforms, such as Windows, Linux, Macs, UNIX, and so on) can affect interoperability with other systems due to different versions of software, different network protocols, security methods, and so on. A diverse environment can also affect supportability because the organization must maintain different skillsets and a wide knowledge base to support the diverse platforms.

On the other hand, maintaining a homogeneous environment can reduce costs, ensure interoperability, and allow a more common set of security controls and mechanisms, such as patch and configuration management. However, there is even risk involved in a

homogeneous platform environment because of the likelihood that a vulnerability discovered in one system would also be shared in many others, offering a wider attack vector for a potential malicious actor. It is really a matter of the systems development lifecycle as to how and when platforms are developed, introduced into the infrastructure, implemented, maintained, and, eventually, disposed of, and there is risk inherent in all of these different phases, as I discuss in more detail later.

Cloud

No book on risk would be complete without considering cloud-based risks. The concept of *Everything-as-a-Service (XaaS)* means that, for a price, nearly anything a user or company would use a computing system for can be delivered to them by a cloud service provider through the cloud infrastructure—typically through a thin client, web interface, or *application programming interface (API)*. In such a *serverless architecture*, a cloud provider manages the server and associated architecture and allocates resources as necessary to the end users. This requires less upfront investment in hardware and software by the customer, greater scalability, and centralized management for administration and security concerns. Each cloud infrastructure is a little different, but there are three main types of cloud infrastructure to consider when looking at risk:

- *Software-as-a-Service (SaaS)* allows a customer to essentially lease software, such as applications and databases, thus enabling rapid rollout to the greater user community.
- *Platform-as-a-Service (PaaS)* provides the framework of an operating system and associated software required to perform a function (e.g., the Linux operating system and components needed to run a web server).
- *Infrastructure-as-a-Service (IaaS)* provides the ability to quickly stand up virtual machines, storage devices, and other infrastructure that would otherwise require the purchase of physical devices.

Project, Program, and Portfolio Management

Organizations use project, program, and portfolio management to oversee and sustain both short- and long-term aspects of systems and processes. These categories apply to the scope and scale of different sets of activities or processes within an organization. *Projects* are limited-duration sets of activities geared toward a particular goal; *programs* are ongoing, longer-term, and may also encompass several individual projects as well as other activities specific to processes that may have an indefinite duration. *Portfolio* management is the oversight of multiple programs by a senior person within the organization. Note that the major difference between a project and a program is duration; a project has a definite beginning and end, whereas a program is indefinite. *Project management* represents the processes and practices used by organizations to manage individual projects.

Project Drivers and Risks

At the core of a project are three primary drivers: the *scope* (the amount or range of work), *schedule* (when work is to be done and its completion date), and *cost* (including all the resources expended toward the completion of the project). All projects depend on these three elements to ensure success. Any shortfalls in cost, delays in schedule, or out-of-scope work (also called *scope creep*) affect the ability of the project to succeed, on time, within budget, and have all the agreed-upon work completed. The following sections discuss the risk factors that can affect all three of these elements.

Scope Risk

Since the three major elements of a project or program are scope, schedule, and cost, risk factors associated with projects and programs primarily affect each of these elements. Factors that affect scope often come from disagreement on exactly what work must be accomplished, in what order, and who will do it. These factors should be examined and agreed upon during the requirements process, where the project is formally scoped. In complex projects, factors that may affect scope include the necessity to task workers from different sections or departments, the exact amount of work to be done, in what order it will be done, and the priority of each task and subtask. The political culture within the organization can adversely affect scope as well, since there may be disagreement between departments or executives on what work is required and whose responsibility it is to accomplish it.

Schedule Risk

The scheduling component of a project has several risk factors as well. The amount of work to be done (scope) affects the schedule as a risk factor because of the time given to accomplish a specific amount of work and the resources allocated to this work. Lack of workers or materials adversely affects the schedule, as does the availability of equipment, work location, or site, and external factors such as contract negotiations, government shutdowns, and so on, as examples.

Cost Risk

Cost risk factors stem from a lack of budget control during the project, including an inaccurate budget or estimation of cost, unexpected expenditures, a weak economy, and so on. Project costs are affected by the cost of labor (to include wages, insurance, and benefits) as well as the cost of supplies or materials needed to complete the project. Additionally, since all three of the elements of scope, schedule, and cost are also risk factors for each other, a delay in schedule or any extra, out-of-scope work impacts the project budget and costs as well.

Project Threats

The threats to a project or program are similar; the major difference is the duration of the time that the threat represents potential harm to a project or program. Keep in mind that threats for a project or program aren't those you might traditionally think of as threat

sources; they are rarely malicious in nature and usually come from the business environment itself. For example, threats to cost are usually those that involve money or resource issues, such as price increases, currency value fluctuation, stock market variations, and so on. These are usually external threats; internal threats to costs might include a sudden budget restructuring, cutting funds for a project, bankruptcy, and so on. Threats to the schedule might consist of worker shortage, strikes, inadequate training, delays in contract negotiations, delays from suppliers, and so on. Scope threats usually involve extra or additional work due to faulty requirements gathering and decisions about what falls into the project's scope and scale.



Sometimes it's difficult to distinguish between threats to each of these elements individually and threats that affect all three project elements simultaneously, as it is more often the case that any threat affecting one element will usually, even indirectly, affect all three.

Project Vulnerabilities

Since vulnerabilities are weaknesses or a lack of protective controls, the vulnerabilities found in project and program management are inherent to weaknesses in their critical elements and processes. For example, a weakness often found in the cost element of a project or program is a failure of the organization to adequately budget for all the resources it requires for the project. The organization might not list every piece of equipment or every major supply it needs, or consider labor costs such as overtime in the event of schedule slippage. It may not negotiate firm prices with a supplier, who may then raise prices on critical parts or equipment. These are all vulnerabilities inherent to the budgeting and cost elements of a project.

Scheduling also has common vulnerabilities. If an organization fails to build in potential scheduling delays, such as work stoppages, dependencies on other organizations to fulfill parts of the schedule, or even holidays, this can lead to problems with getting the project completed on time. Another scheduling vulnerability is to incorrectly estimate the amount of time it takes to perform a task or get a piece of work accomplished.

Vulnerabilities inherent to the scope element of a project affect how much work must be done, in what order, and by whom. A *work breakdown structure (WBS)* is a document that the project manager develops that covers all this information and more. The WBS describes, in excruciating detail, the work, its subcomponent parts and tasks, what skillsets are needed, and what resources are required to do the job. Other documents may break down individual tasks in greater detail, even describing step-by-step how a task is accomplished. If the project manager and team do not create these documents, this is a weakness, as the work may not be performed exactly as needed. Resources may not be adequately allocated for the work, or the work may be performed by unskilled persons. Additionally, the amount and quality of work may not meet the requirements. Failure to set requirements at the outset of the project and to periodically check and track the work against those requirements is also a vulnerability, as is scope-related vulnerability, since it can affect whether the work meets the established requirements.

Project Risk Monitoring

An additional vulnerability that simultaneously affects all three project elements is the failure to monitor or control all these elements on an ongoing basis throughout the life of the project or program. Scope, schedule, and cost must be monitored continuously to ensure they meet the original requirements outlined in the project's charter. The project or program manager must monitor and control these things to prevent cost overruns, schedule slippage, and scope creep. Failure to monitor or control these elements is a serious vulnerability and could jeopardize the success of the entire project or program. Table 4.2 presents examples of risk factors, threats, and vulnerabilities related to project and program management.

Other Project Risks

Other risk factors for projects include training, staffing the project with the right mixture of personnel who have the right skillsets, attitudes, and focus on making the project or program a success. A manager without the right project management skillset is a risk factor that can affect scope, schedule, and cost if they do not manage those three critical elements properly. A technician without the right technical skills may cost the project more money in training or delay work unnecessarily due to rework or slow production.

In addition to factors such as scope, cost, and schedule, adding an indefinite or longer-term duration affects programs as a risk factor because you must project how they will be staffed, funded, equipped, and otherwise managed in the future, over a more

TABLE 4.2 Examples of project risk factors, threats, and vulnerabilities.

Risk Factor	Threat	Threat Agent	Vulnerability	Affects
Unstable economy	Cost increases for critical parts	Supplier	Tight budget or inexact budget estimates	Cost
Inability to accurately define scope during requirements meetings	Scope creep; inexact work specifications	Different department managers and supervisors	Lack of defined scope or work specifications	Scope
Manufacturing expensive parts for sensitive components	Delays due to inability to receive spare or replacement parts	Supplier	Lack of critical parts	Schedule
Training levels/ skillsets	Incomplete or shoddy work	Untrained worker	Precision or delicate work required	All

extended, indeterminate period. Therefore, risk factors could be long-term and fluctuate over time. Factors such as market, economy, and technology changes will influence risk in a program much more so than a project, affecting its scope, schedule, and cost over a more extended period.

Modern Development Models and Risks

The first several decades of system and software development relied on the time-proven *waterfall* methodology. This methodology was thorough, but slow. As a result, the IT industry has developed methodologies that enable systems to be developed and implemented more quickly without sacrificing quality. These methods include Agile, DevOps, DevSecOps, CI/CD, and Infrastructure-as-Code (IaC).

Modern system delivery models emphasize speed, adaptability, risk reduction, and incremental value. These approaches reduce the long planning cycles of traditional waterfall methods but introduce distinct risks related to the pace of change, automation, tooling complexity, and reduced documentation. Risk practitioners must understand how these models alter control expectations, testing practices, and governance requirements.

Agile

Agile delivery uses iterative development cycles, known as *sprints*, to deliver incremental functionality, whereas the Waterfall model delivers complete functionality all at once. Frequent releases and evolving requirements reduce time-to-value, but create risks related to incomplete documentation, inadequate long-term planning, and inconsistent control implementation. Rapid iterations may limit formal design reviews, expand technical debt, and allow defects to propagate across sprints. Risk management activities must be embedded into sprint planning, definition of done, and backlog refinement to ensure traceability and control coverage.

Scrum

Scrum structures Agile delivery using defined roles (Product Owner, Scrum Master, Development Team) and events (sprints, reviews, retrospectives). Risks arise when product backlogs lack prioritization, sprint goals are unclear, or teams bypass retrospective-driven improvements. Dependencies among teams can produce integration failures if coordination mechanisms, version alignment, and interface testing are insufficient. Weak sprint governance may also result in scope instability or inconsistent quality across increments.

DevOps/DevSecOps Pipelines

DevOps integrates development and operations to accelerate delivery, relying on automation, continuous testing, and shared responsibility. *DevSecOps* extends this model by embedding security controls directly into the pipeline. Risks include insecure automation scripts, poor secret management, lack of segregation between build and production environments, and insufficient visibility into automated changes. Pipeline failures or compromised tooling can directly affect production systems due to high degrees of automation and elevated privileges.

CI/CD, Infrastructure-as-Code

Continuous Integration/Continuous Delivery (CI/CD) automates code integration, testing, packaging, and deployment. *Infrastructure-as-Code (IaC)* applies version-controlled code to manage infrastructure. These practices reduce manual effort but introduce risks when automated tests are incomplete, deployment gates are misconfigured, or IaC templates propagate insecure configurations. Mismanaged pipelines may deploy untested or vulnerable code, while configuration drift can occur if manual changes bypass IaC processes. The integrity of build artifacts and pipeline components is critical to operational security.

Disaster Recovery Management

Business continuity and disaster recovery management are essential parts of operating a business, regardless of its market context. This is because incidents, disasters, and other disruptive events happen to all organizations eventually, despite the most favorable conditions and the most careful planning. Business continuity and disaster recovery primarily aim to achieve business *resilience*.

Disaster recovery planning (DRP) is part of the broader discipline of business continuity planning (BCP), which focuses on measures to ensure the continuity of critical business processes and the organization's survival. To retain alignment with the CRISC job practice, BCP is covered in Chapter 5.

DRP concerns measures for recovering individual IT components when a disruption occurs. In response to a disaster, disaster recovery plans include procedures for the safety of facilities and personnel, such as evacuation and first aid.

DRP is driven by *IT service continuity management (ITSCM)*, which ensures that critical IT services can continue or be restored after a disruption. Figure 4.2 depicts the relationship between these disciplines. ITSCM focuses on service-level performance, whereas DRP operates at the component level. ITSCM and disaster recovery are closely related disciplines that operate at different layers of continuity planning. Table 4.3 highlights their distinction.

FIGURE 4.2 The relationship among BCP, ITSCM, and DRP.

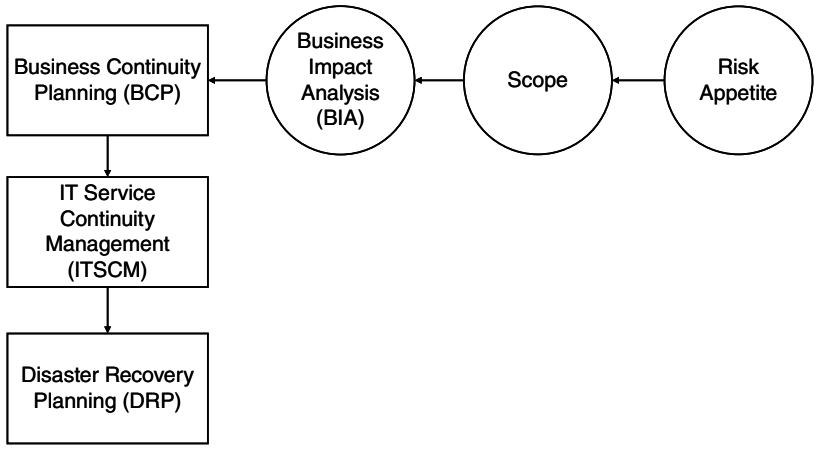


TABLE 4.3 ITSCM and DR comparison.

Dimension	IT Service Continuity Management (ITSCM)	Disaster Recovery (DR)	Relationship/Mapping
Primary Purpose	Ensure that critical IT services can continue following a major disruption (or be restored to an acceptable level).	Restore technology components (systems, applications, infrastructure, data, facilities) to operational status after a disruption.	ITSCM defines required service levels; DR provides the technical means to achieve them.
Scope	End-to-end IT services, including applications, dependencies, processes, service levels, and business-facing outcomes.	Technical infrastructure: servers, storage, networks, cloud environments, identity services, data recovery mechanisms.	ITSCM = service-centric. DR = component-centric.
Drivers	Business Impact Analysis (BIA), service-level requirements, risk appetite, resilience objectives.	RTO, RPO, technical recovery capabilities, infrastructure recovery strategies.	ITSCM translates BIA outputs into service requirements; DR translates these into system recovery mechanisms.
Typical Activities	• Define IT service continuity requirements • Develop end-to-end continuity architectures • Ensure provider capabilities match service needs • Validate continuity through service-level testing	• Execute backup/restore • Failover/failback • Rebuild infrastructure • Perform technical DR testing	DR is one functional component executed within ITSCM.
Planning Artifacts	• ITSCM strategy • Service continuity plans • Continuity architecture diagrams • IT service resilience requirements	• DR plan/runbooks • Recovery architecture • Backup strategies, replication configs • DR test documentation	ITSCM artifacts define "what must be continuous." DR artifacts define "how systems recover."

(*Continued*)

TABLE 4.3 (Continued)

Dimension	IT Service Continuity Management (ITSCM)	Disaster Recovery (DR)	Relationship/Mapping
Ownership	IT service management leadership; continuity managers; service owners.	Infrastructure, operations, and platform teams.	ITSCM sets requirements; DR implements them.
Success Criteria	Ability to maintain or rapidly restore IT services at agreed levels.	Ability to restore technical systems within defined RTO/RPO.	Service continuity success depends on DR success.
Risk Focus	Service-level risk (service outage duration, cross-dependency failures, provider limitations).	Technology-level risk (data loss, system rebuild failure, backup integrity, infrastructure availability).	ITSCM manages service risk; DR manages technical recovery risk.
Testing	End-to-end service continuity exercises; validation of service dependencies.	Technical DR tests (tabletop, simulation, parallel, full interruption).	ITSCM uses DR test results to assess service continuity capability.



In the CRISC job practice, ISACA categorizes disaster recovery management as an IT activity and business continuity management as an information security activity. Thus, disaster recovery management is discussed in this chapter, while business continuity management is discussed in Chapter 5. Despite this logical separation, the two should ultimately be managed as a single mission: ensuring business survival in the face of significant threat events.

DR Concepts and Planning Activities

Disaster recovery (DR) focuses on restoring IT services, data, and supporting infrastructure following a disruptive event. Effective DR planning ensures that critical systems can resume operation within acceptable time and data-loss thresholds defined by business continuity requirements. DR activities emphasize preparedness, structured response, and controlled restoration of technology assets and the facilities where they reside.

Effective DR planning requires alignment with business continuity objectives, disciplined documentation, accurate dependency mapping, and a structured approach to restoring critical technology services following disruptive events.



Not All Agree on BCP and DRP Boundaries

Ask ten information security professionals for the definition and scope of BCP and DRP, and you're likely to get ten different answers. For some, they are synonymous; for others, they are related; and for yet others, they are unrelated.

This text aligns with the traditional definitions: business continuity emphasizes the resilience of business processes and procedures, while disaster recovery emphasizes the resilience of supporting information systems. They work together to ensure organizational continuity during and after a disruptive event.

DR Objectives and Scope

DR planning begins with defining the systems, applications, and infrastructure components covered by the DR program. Objectives must align with the organization's *business impact analysis (BIA)*, which determines mission-critical services and their required availability. Scope considerations include recovery of data centers, cloud workloads, networks, identity services, and essential third-party dependencies. BIA is explored in more detail in Chapter 5.

Recovery Time Objective (RTO) and Recovery Point Objective (RPO)

Driven by the BIA, *recovery time objective (RTO)* specifies the maximum acceptable time to restore a system after an outage. *Recovery point objective (RPO)* defines the maximum tolerable data loss measured in time. These parameters guide decisions on backup frequency,

backup media management, replication strategy, recovery site type, and technology investments. Systems with low RTO or RPO values require automation, near-real-time synchronization, and highly resilient architectures.

Related recovery objectives include:

- ***Recovery capacity objective (RCapO)*:** This refers to the transaction processing capacity of an alternate processing system that would be pressed into service when the primary system is incapacitated. Capacity can be expensive, and outages are usually brief; senior management must decide on how much to invest in a recovery system that may be seldom used. RCapO is expressed as a percentage: for example, the RCapO of a recovery system can process transactions at 75% of the capacity of the primary system.
- ***Recovery consistency objective (RCO)*:** This refers to the completeness or consistency of features and capabilities of a recovery system as compared to the primary system. Some organizations opt to have their recovery systems include only the core features of their primary systems, to reduce the cost and complexity of the recovery system.

Business Impact Analysis Integration

The BIA identifies critical processes, required resources, upstream and downstream dependencies, and the consequences of extended outages. DR planning uses BIA outputs to prioritize recovery sequences, determine acceptable downtime, and allocate resources to the most critical services. Dependencies across on-premises and cloud workloads are documented to prevent recovery gaps.

Role of DRM in Business Continuity

Disaster Recovery Management (DRM) supports business continuity by restoring the technology components and data required for critical business functions (defined and driven by ITSCM) to resume following a disruptive event. While business continuity defines the broader organizational capability to maintain essential operations, DRM provides the technical means for recovering systems, applications, infrastructure, and connectivity within defined recovery objectives.

DRM ensures that recovery strategies, backup mechanisms, failover procedures, and restoration workflows align with business continuity requirements, including RTO, RPO, and process prioritization. It also supplies the technical coordination needed during incident response, enabling continuity teams to activate contingency procedures while IT restores supporting services.

Through structured planning, documentation, testing, and continual improvement, DRM forms the technical foundation that enables the organization to meet its continuity commitments and maintain resilience during and after a disruption.

Risk Assessment and Scenario Identification

Risk assessments identify events that could trigger DR activation, such as natural disasters, cyberattacks, infrastructure failures, or extended outages at service providers. Scenario analysis helps determine which recovery strategies are appropriate for different types of disruptions. DR plans must consider loss of facilities, network unavailability, data corruption, identity service failures, and ransomware impacts.



DR-related risk assessments should also consider the possibility that DR plans may not function as intended, potentially extending the length and impact of outages.

Plan Documentation

The DR plan documents roles, responsibilities, recovery procedures, communication protocols, escalation paths, and step-by-step restoration activities. Documentation must be accurate, accessible during a disaster, version-controlled, and stored in multiple secure locations. Plans should include contact lists, vendor support information, recovery topology diagrams, and validation checklists.



DR plans should be reviewed at least annually and whenever significant changes are made to systems. Arguably, an organization's IT change management process should include changes made to DR plans, so that they always remain in sync with system recovery procedures.

DR Training and Preparedness

Personnel must be trained in their DR responsibilities, including technical tasks, command-and-control roles, and communication requirements. Training improves readiness and reduces errors during high-stress recovery operations. Regular exercises reinforce procedures and identify weaknesses in the plan.

Disaster Declaration, Activation Procedures

Disaster declaration and activation procedures define when the disaster recovery plan is initiated and how the organization transitions from normal operations to recovery mode. Clear criteria and decision-making authority are essential to avoid delays, confusion, or inconsistent responses during disruptive events.

A disaster is declared when an incident exceeds the organization's ability to restore services through standard operational processes or when recovery time requirements cannot be met without activating DR capabilities. Common triggers include sustained system outages, data center failures, cybersecurity incidents, loss of critical infrastructure, or prolonged service provider disruptions.

The activation process identifies authorized personnel—typically senior IT, continuity, or incident management leaders—who assess the situation, confirm the severity, and formally initiate the DR plan. Activation procedures include notifying recovery teams, initiating communication protocols, mobilizing resources, accessing recovery sites or cloud environments, and executing documented restoration steps.

Clear, repeatable activation procedures ensure timely engagement of recovery processes and provide a controlled transition into full disaster recovery operations.

DR Strategies

Disaster recovery strategies define the technical and procedural methods an organization will use to restore systems, data, and services following a disruptive event. Strategies translate recovery requirements—such as RTO, RPO, resource availability, and dependency priorities—into actionable restoration approaches. These strategies must align with the organization's business impact analysis, continuity objectives, and available recovery capabilities.

The recovery strategy developed by the organization is a result of careful planning, as well as budgeting and allocating resources to the continuity management efforts. The recovery strategy should address the most efficient methods of getting the business back online and functioning after properly responding to the critical needs of protecting people and equipment during the disaster itself. Keep in mind that recovering the business to an operational state does not necessarily mean it will be working at 100% of its previous operational capacity; the business may only be able to get a few key services up immediately following a disaster. The organization should take a realistic look at the degree of continuity possible, given the seriousness and scope of the disaster, its resources, the environment it is working with after the disaster, and the availability of people, supplies, infrastructure, facilities, and equipment. All of this should be considered in the recovery strategy planning.

DR Strategy Development

A DR strategy defines how each system will be restored. Strategies may include backup-based recovery, failover to alternate sites, cloud-based replication, or rebuild-from-code approaches for Infrastructure-as-Code environments. Strategy development includes identifying the minimum resources required for recovery, such as compute capacity, storage, personnel, and access to critical documentation.

Based on the priority of systems as defined by the BIA, a DR strategy will also implicitly define which less-critical IT systems will be excluded from DR planning.

Backup-based Recovery

Backup-based recovery relies on stored copies of data and system images on backup media, including disk, tape, cloud backup repositories, or immutable storage.

Key risks and considerations include:

- Alignment of backup frequency to RPO requirements.
- Verification of backup integrity through routine testing.
- Secure storage and access controls to prevent tampering or ransomware encryption.
- Documented procedures for restoring systems, applications, and configurations.

Backup-based recovery is suitable for systems with moderate RTOs or limited interdependencies but may be insufficient for mission-critical workloads requiring rapid restoration.

Redundant or High-availability Architectures

Redundant architectures use fault-tolerant systems, clustering, geo-clustering, load balancing, and other high-availability configurations to maintain service continuity during localized failures.

Benefits include minimal downtime and reduced reliance on manual recovery. However, redundancy does not replace DR; failures such as data corruption, ransomware, or major site outages still require full recovery procedures.

Replication and Mirroring

Data replication strategies synchronize data between a primary environment and one or more secondary sites. Options include synchronous replication (zero or near-zero RPO) and asynchronous replication (RPO measured in seconds or minutes).

Replication supports rapid recovery but requires:

- Sufficient bandwidth and storage capacity
- Compatibility between primary and secondary environments
- Controls to prevent replicated corruption or malware propagation

This strategy is used for critical systems with low tolerance for data loss.

Alternate Processing Sites

DR strategies may include the use of hot, warm, or cold sites (or cloud-based equivalents) to restore operations:

- **Hot sites** provide fully operational systems with near-real-time data.
- **Warm sites** maintain partially configured infrastructure requiring some setup.
- **Cold sites** provide space and utilities but require full system rebuilds.

Proper selection depends on RTO, resource availability, and cost constraints.

Cloud-based DR strategies use public or hybrid cloud environments for backup storage, system replicas, or failover capacity. Techniques include:

- Backup-to-cloud with on-demand restoration
- Cloud-based standby systems
- Cross-region or cross-availability-zone failover
- Infrastructure-as-Code (IaC) rebuilds in cloud environments

Cloud approaches reduce capital costs but require careful configuration, identity alignment, and validation of cloud-provider SLAs.

Rebuild-from-code and Automation

For modern architectures, particularly containerized or Infrastructure-as-Code (IaC) environments, systems can be rebuilt automatically from stored configuration templates, deployment scripts, and container images. This reduces reliance on physical infrastructure and speeds recovery, but requires disciplined configuration management and secure code repositories.

Prioritized Recovery Sequencing

A complete DR strategy defines the order in which systems will be restored. Sequencing accounts for application dependencies, identity services, networking, storage, and integration points. Typical sequences restore foundational services first, followed by critical business systems, and then lower-priority workloads.

Sequencing should be driven by BIA results, RPO and RTO targets, and service dependencies. For example, while identity and directory service applications are not themselves high business priorities, it may be necessary to recover those before critical business applications are recovered.

DR Plan Testing

Testing the disaster recovery plan is critical; without testing, the organization will not know whether it's effective until a disaster occurs. Normally, that's a bit too late to find out that you haven't carefully considered all the different events that could occur during a disaster and how you will deal with them. Testing the plan involves making sure that people know what their responsibilities are, that they are well-trained, that they have all the equipment and resources needed to adequately respond to a disaster, and that they practice these responses. Testing the disaster recovery plan will point out deficiencies in training, equipment, resources, and response activities.

Types of DR Testing

Multiple test types provide increasing levels of rigor and realism.

- **Readthrough test:** Individual reviews by recovery team members and others of recovery procedures. There is no discussion among team members, as reviews are performed individually.
- **Walkthrough (tabletop) test:** A discussion-based exercise where recovery team members review the DR plan, roles, dependencies, and decision points. No systems are activated. This identifies documentation gaps, unclear procedures, and coordination issues.
- **Communication test:** A focused test of notification procedures, contact lists, escalation paths, and communication channels. Ensures that personnel and external parties can be reached quickly and that communication tools function during an incident. Note that communications to external parties are *not* sent.
- **Checklist/documentation test:** Teams verify the accuracy, completeness, and accessibility of recovery documentation, inventories, runbooks, and configuration data. No technical restoration occurs.
- **Simulation test:** A structured scenario is executed without interrupting production, simulating failures or outages. Technical teams may use non-production systems to practice restoration steps.
- **Parallel test:** Recovery systems are brought online in a separate environment while production continues running. This validates technical recovery steps, data integrity, and system functionality without risk to live services.

- **Full interruption test:** Production systems are intentionally shut down, and recovery operations are executed fully. This provides the highest assurance but carries the greatest operational risk and is used sparingly, typically under controlled conditions.

Effective programs use a mix of test types to progressively validate capabilities while minimizing operational disruption.



Because some types of disasters result in recovery team members being unable to respond, recovery documentation should be written for audiences less familiar with the organization's environment. Testing should also include additional personnel who may be called upon to respond during an actual disaster.

DR Testing Metrics

DR testing metrics assess the effectiveness, reliability, and completeness of recovery processes. Metrics support governance, identify weaknesses, and validate alignment with business continuity requirements.

- **Recovery Time Objective (RTO) achievement:** Measures whether systems were restored within the allowed downtime. Variances indicate process inefficiencies, resource deficits, or unrealistic planning assumptions.
- **Recovery Point Objective (RPO) achievement:** Evaluates whether data loss stayed within acceptable limits. Failures may indicate inadequate backup frequency, replication gaps, or delayed synchronization.
- ***Mean Time to Recover (MTTR)* achievement:** The actual time required to restore a system during testing. MTTR provides a realistic estimate of operational recovery capability and highlights bottlenecks.
- **Procedure accuracy and completeness:** Tracks the number of steps that were unclear, incorrect, or missing from DR documentation. High variance indicates a need for updates or retraining.
- **Resource readiness:** Assesses availability and adequacy of personnel, tools, backup media, access credentials, and infrastructure required for recovery.
- **Test success rate:** Measures the percentage of test objectives achieved across all planned recovery activities. Partial or failed results highlight control or procedural deficiencies.
- **Issue identification and remediation rate:** Tracks the number of findings generated during tests and the time required to resolve them. Demonstrates the maturity of the DR improvement process.
- **System dependency validation:** Evaluates whether all required upstream and downstream dependencies were identified and correctly restored in sequence.

Consistent measurement of these metrics supports continuous improvement and ensures the DR program maintains alignment with evolving business and technology requirements.



Much more information on business continuity and disaster recovery planning can be found in two standards:

- National Institute of Standards and Technology (NIST) Special Publication 800-34, "Contingency Planning Guide for Federal Information Systems," which was written for use by government agencies but can be used by anyone or any business to help in understanding and planning for business continuity.
- ISO 22301, "Security and resilience—Business continuity management systems—Requirements," an international standard for business continuity management.

Data Lifecycle Management

Most organizations produce a wide variety of information, from procedure documentation consisting of publications for internal use, to confidential papers for senior management, to publicly available documents, as well as business records stored in file shares and in business applications' databases. Without proper controls, such information could be used to compromise an organization's security. Every organization should have guidelines that ensure all information produced by the organization is classified, organized, and stored securely to prevent its loss, damage, unauthorized modification, or theft.

Data Governance

As part of its governance and risk management programs, an organization should develop a data governance program that includes data classification and handling policies and lifecycle management controls. Relevant activities and controls include:

- **Data inventory:** Like other types of tangible and intangible assets, organizations track their inventory of data, including various types of metadata implied by the subsequent items in this list. Like other types of asset management, organizations need to ensure that their data inventories are complete and accurate.
- **Data ownership and stewardship:** Policy that specifies the position, title, or role of the person responsible for each database and dataset, together with specific responsibilities, such as access requests, change requests, and data quality.
- **Data classification:** A policy that defines security levels with respective labeling and handling procedures, discussed later in this subsection.
- ***Data provenance*:** Records that describe the origin of data in its possession, whether obtained from customers, data brokers, or created in-house.
- ***Data lineage*:** Records that describe post-creation/acquisition additions, changes, and deletions of data.

- ***Data minimization*:** The practice of receiving and storing only those data records and data fields necessary to perform business operations and provide services. Receiving more than what is necessary increases risk if data is stolen by cybercriminals.
- **Data retention:** A policy that defines the minimum and maximum period of time that specific databases and datasets should be retained. An exception process is required to accommodate situations in which data is needed beyond its retention period.
- ***Legal holds*:** An exception process related to data retention in which an organization's legal department may advise IT that specific data may be needed in a legal discovery process, and that all such data should be exempt from modification and deletion until further notice.
- **Data access:** Policy and controls that stipulate the persons, groups, and roles, with specific access rights to databases and datasets.
- **Regulatory constraints:** Laws, regulations, and contractual stipulations that govern the retention, use, and transfer of data. Examples include privacy, cybersecurity, and AI-related regulations.

Data Ownership and Stewardship

Data governance and data management include activities related to the ongoing management of databases and datasets, including:

- **Data classification:** The classification level assigned to each database and dataset. This is explained in more detail in the next section.
- **Access requests:** For each request to access data (whether reading, writing, modifying, or removal), a designated data owner will review and grant or deny access to data.
- **Access reviews:** For each database and dataset, owners and stewards should periodically review the persons, groups, and roles that have access, to determine whether such access needs to continue.
- **Data quality:** Owners and stewards may be involved in a variety of activities to ensure a minimum level of data quality and completeness.

Data Classification

To aid in the document management effort, organizations implement a *data classification* policy, where documents, databases, and datasets need to be assigned security classifications to indicate their level of confidentiality and then labeled appropriately. Each classification requires different standards and procedures of access, distribution, and storage. The classification also sets a minimum standard of privileges required by a user to access that data. If a user doesn't have the necessary access privileges for that classification of data, the user won't be able to access it. Typically, access is delineated using subjective levels such as high, medium, and low. These should be agreed upon by management based on

the sensitivity of the data and the damage to the organization if the data is subjected to unauthorized access.

Several levels of classification can be assigned, depending on the type of organization and its activities. A typical organization might have only two classifications: private and public. *Private classified documents* are intended only for the internal use of the organization and cannot be distributed to anyone outside the organization. *Public documents*, however, would be available to anyone.

Government and military institutions might have several levels of confidentiality, such as Unclassified, Confidential, Secret, Top Secret, and so on. Each level of classification represents the level of severity if that information is leaked. For example, the lowest level (Unclassified) means that the document is not considered confidential or damaging to security and can be more freely distributed (though not necessarily releasable publicly). At the highest level (Top Secret), documents are highly restricted and would be severely damaging to national security if they fell into the wrong hands. Each document needs to be assigned a classification depending on the sensitivity of its data, its value to the organization, its value to other organizations in terms of business competition, the importance of its integrity, and the legal aspects of storing and distributing that data.

Implementing Data Classification Policy

The cybersecurity leader in a high-tech instruments manufacturer has decided to implement a data classification program. Because the organization has not had a data classification program in the past, the effort required considerable planning, and the organization estimated that the endeavor would take three to five years to complete, including automation.

The cybersecurity leader leveraged relationships with company departments and spoke with each department head about the initiative. All agreed that this effort would begin with some workforce messaging to announce the initiative. Next, the data classification policy was drafted and circulated for comment, and further workforce messaging was provided to describe data classification and what is expected of them.

The cybersecurity leader selected two "friendly" departments, HR and Legal, to begin working with those teams to ensure that team members understood data classification and the impact it might have on their work. Then, as each additional department was added, the workforce gradually became accustomed to data classification. Finally, the company began implementing automation, particularly a *data loss prevention (DLP)* system that is integrated with email and file servers.



The security protections, access controls, data retention, and storage and disposal policies to be used depend on a document's security classification.

The Data Lifecycle

As a part of their business operations, organizations create, acquire, manage, manipulate, share, and eventually delete data. These and other activities can be described in a lifecycle process that includes:

- **Creation, ingestion, or acquisition:** The origin of data, whether created within an organization's information systems, or acquired from an external entity such as a data broker or from individual customers.
- **Storage:** Once acquired, data is stored in a database or dataset that resides in a physical system such as a database server or a file server. The criticality of stored data, together with recovery targets defined in a business impact analysis (BIA), will influence the required resilience, including replication, mirroring, and backup. Data classification policy may define where databases and datasets must be stored, and whether they must be encrypted in storage.
- **Access:** Stored data is accessed as a part of one or more business processes. A suite of access controls determines the persons, groups, roles, and systems with permission to access data. Such accesses are often logged in an audit log that can be used later to determine who or what accessed data at some point in the past.
- **Transmission:** Stored data may be transmitted to various systems within an organization, and may also be shared with external parties as part of business processes. Data classification policy may permit and prohibit forms of transmission and determine whether data must be encrypted in transit.
- **Processing:** Data governance and data classification may stipulate how databases and datasets may be processed, to ensure that each database and dataset is processed only by systems and in a manner approved by management.
- **Sharing:** Organizations may utilize third-party service providers to perform some processing of data. Policy, governance, and third-party risk management will all influence how this sharing is controlled and managed.
- **Retention:** Organizations may develop a data retention schedule, which stipulates the minimum and maximum periods of time that data records must be retained. Retention is usually driven by regulation, as well as risk.
- ***Archival*:** Archival is the long-term storage of data, often to comply with data retention regulations. The data architecture of some information systems may permit portions of their databases or datasets to be removed from an active business application and moved to a long-term storage system. Archiving may also involve storing data at an alternate storage facility.
- **Deletion:** When data records, or entire databases and datasets, reach the end of their stated retention, such data is removed from information systems. The nature of the systems and media storing this data, and the nature of the data itself, will contribute to the method(s) used to delete data. These methods range from simple deletion to military-grade overwriting, as well as physical destruction of storage media. Organizations often develop *data destruction* procedures to stipulate required methods for destroying data, whether it resides as hardcopy, or in a database, dataset, filesystem, backup media, unused PC or server, copier/scanner, or cloud environment.

Systems Development Lifecycle

The *systems development lifecycle* (sometimes known as the SDLC, which, confusingly, is used as the acronym for software development as well) is the entire lifecycle of a system, product, or piece of software. The lifecycle covers every single aspect of a system or product, from the concept to development, from implementation to disposal. There are several popular systems development lifecycle models, including waterfall, iterative, and Agile. I don't cover any of these models here, as an in-depth discussion on any one of them would fill an entire book. However, I do mention some of the more generic parts or phases of the SDLC and briefly describe them so that we can put them in the context of risk factors, threats, and vulnerabilities.



Some models are linear, in that they progress in a sequential manner only, and some are cyclic in some way (repetitive, recursive, or iterative), where a system can be reintroduced into an earlier phase for updates or upgrades or to fix problems or issues.

SDLC Phases

At the basic level, the SDLC consists of several phases that include conceptualizing the system, discovering the requirements for a system, designing its architecture, developing the system, testing it as both a single unit and as integrated with larger systems in its environment, implementing the system, and, finally, disposing of the system after it has reached the end of its useful life. Each of these different phases of the lifecycle has different aspects that include planning, management, resourcing, and so on. Each also has inherent risks. The next few subsections discuss these phases generically.



Some of the names given to the different lifecycle phases may vary among different lifecycle models, but in general, the concepts are the same.

Planning

The first phase of the SDLC normally involves some type of conceptual planning process, where the need for the system is determined, and a general idea is developed of what the system needs to do and what purpose it will serve. Some models refer to this as the *initiation phase*. During this phase, the feasibility of developing or acquiring a system is studied. It includes considerations such as when the system is needed, how much it might cost, what type of staffing it might need to develop, implement, and maintain it, what other systems it might need to connect to or interface with, and so on. Security and privacy professionals want to be involved in this phase of the lifecycle so that they can understand the basic concepts of the system and provide feedback on any potential security or privacy risks.

Often, it is the planning phase where a business case is developed, to help the organization determine whether the system (or changes to an existing system) is justifiable. A *business case* describes the resources required to build (or change) the system and the economic benefits expected to be derived from the system.

Requirements

The next phase of the SDLC is generally the *requirements phase*, and this phase is typically named as such in most models. During this phase, different sets of requirements are developed. These requirements might include functional requirements, performance requirements, security requirements, and business requirements. Requirements dictate concrete terms, such as what a system is exactly supposed to do, how it will do it, to what degree, and what standards it must meet. Without a clear set of requirements, there can't be any good traceability back to what the system was supposed to do in the first place. That's why it's very important to try to accurately cover all the possible requirements and needs the system may have to meet. This is also where the scope of the development process is established. Security and privacy professionals should be invited to develop security and privacy requirements for the system, to ensure that later phases of the project will include all required characteristics.

The idea of SMART requirements helps organizations improve the clarity and quality of their requirements. SMART stands for:

- **Specific:** Clearly states the required behavior or constraint without ambiguity.
- **Measurable:** Includes objective criteria for verification (metrics, thresholds, pass/fail conditions).
- **Achievable:** Is technically and operationally feasible within constraints.
- **Relevant:** Directly supports a business objective, risk response, or compliance need.
- **Time-bound:** Specifies timing, frequency, or deadlines where applicable.

SMART requirements are not a separate category of requirements; rather, SMART should reflect attributes of every requirement.



The entire corpus of requirements should be utilized in the testing phase, so that testing is able to verify whether each and every requirement has been correctly implemented.

Design

The next few phases can be a bit different based on the model under discussion and whether the product being developed is a system or software. System developers would move from the requirements phase into designing the architecture in terms of developing a general high-level design of how the system is supposed to work, what major components it might have, and what other systems it might interface with. Software developers, similarly, might design software architecture at a higher level during this phase. As the

SDLC development process moves forward, this overall architecture is broken into smaller pieces until finally individual components or pieces of code can be developed. Again, how this design process progresses depends on which model is being used. Security and privacy professionals should be asked to review the design of the system to ensure that requirements were properly represented.

Development

After the design phase, decisions must be made as to whether to build a system or software or simply acquire it from a third party. Some systems and software can be bought from other developers, but some must be developed specifically for the end application or end use. If the system is to be acquired from some other source, and the design specifications are provided to the best source, then a system or software that meets the requirements and design will be provided. If the system is to be developed, the next phase of the SDLC comes into play. During this phase, the software is written, and systems are assembled from individual components. Organizations can use tools to examine a system's source and object code to identify any security-related defects that could be exploited by an attacker.

Testing

Following the development phase, many SDLC models include a test phase. In the test phase, different aspects of the software and systems are tested. Many different tests can take place during this phase. Some of these tests might involve *unit testing*, which simply means testing an individual component or unit of software at its most basic level for functionality and performance. *Integration testing*, another type of testing, usually involves testing the overall function and performance of all the components of the system or piece of software working together, as they would when assembled. During this phase, other types of testing could take place, including *interface testing* (testing different aspects of the system when it interfaces with other separate systems in its environment) and *security testing* (which may include compliance testing, vulnerability testing, and even penetration testing). *User acceptance testing (UAT)* involves end users of the system performing various functions to determine whether they will accept the system in its current, developed state.

Test plans are generally derived directly from the body of requirements established and finalized earlier in the project. This is why it's important that requirements be observable and measurable.

Implementation

Once a system has been tested and accepted, the implementation phase usually follows. In this phase, systems or software are put into production and used for their originally intended purpose. Some development models describe an overall implementation phase that includes the day-to-day maintenance and operations of the system. Other models move from the implementation phase to a more long-term maintenance or sustainment phase for the product. Regardless of how the phase is described in the model, during these phases, the system is used, undergoes minor maintenance and upgrades, and is periodically tested to make sure it still meets its original design and requirements.

Operations and Maintenance

Once a system has been placed into operation, it will undergo periodic changes, mainly in response to feature requests, incidents, and outages. These changes will all be managed through the defect tracking and remediation process, which involves:

- Maintaining a log of feature request changes and defects
- Assigning feature requests and defects to developers for remediation
- All necessary testing and approvals
- Promotion of fixes to the production environment through the change management process
- Updates to CMDB and other business records

Post-implementation Review

After a system has been implemented and released into a production environment for business use, organizations often perform a post-implementation review, to identify and discuss:

- What went well during the entire SDLC project
- Issues that occurred during the project
- Potential improvements: A) to the system that was implemented, and B) to the SDLC process
- Action items to be performed by designated individuals

Retirement

Once a system has been in operation for a while, it may eventually be superseded or replaced by a more efficient or updated system. When the system is replaced or removed from the environment, it is said to have entered the *retirement phase*. Some models also refer to this as the *disposal phase*. During this phase, the system is taken out of production, dismantled, and no longer used. Its components may be reused, or they may be disposed of.

Often, when organizations retire a system, a new system is implemented to continue supporting business processes when the organization decides that the old system should no longer be maintained. In terms of dealing with historic data, options for continued business process support include:

- **Migrate all historic data:** To keep the business environment simple, some organizations will migrate all historic data to a new system. Generally, this approach is taken when organizations are required to retain transaction data for many years. In this approach, the old system can be shut down entirely soon after the migration.
- **Migrate no historic data:** To simplify the new system, all historic transaction information is retained on the system being retired. In this instance, the retired system may be modified so that it operates in a "read-only" or "reporting-only" mode to satisfy data retention requirements.

- **Migrate only recent historic data:** This is the compromise between the first two options, where an organization might, for instance, opt to migrate only the past two years of transaction data, as older transactions are rarely needed. Another option is to use the old system to create reports representing all historic data before shutting it down, and then store those reports where staff can find them as required.



While you will likely not be expected to know the details of any specific SDLC model, be familiar with the generic phases of the SDLC: Initiation, Requirements, Design, Development, Testing, Implementation, and Retirement. Also, keep in mind that different models may refer to these phases by different names or even combine or separate phases to some degree.

The Roots of SDLC

It should be noted that in most organizations, the skills to properly set up computers and develop software are completely separate from one another. In fact, many compliance frameworks consider it a serious risk to have the software coders (often called developers) have access to a live production environment, some focusing on the ability of the software coders to implement their code in production. This need to prevent this is called *separation of duties* (sometimes called *segregation of duties*). This came down from the banking industry in that two people were required to perform a task, which means that the likelihood of abuse is drastically reduced. This is another critical factor for risk assessors to consider.

As a result of the differentiation, a *software development lifecycle (SDLC)* is often used as a separate term from systems development lifecycle—two completely different groups and sets of governance are part of the organization. Typically, the CIO oversees an IT organization, while the CTO oversees the software development within an organization. As a result, for many organizations, SDLC typically refers to the former and not the latter interpretation. Both are extremely critical for many organizations, but the approach is different, even though the goals may be the same.

Another factor today would be that with many cloud technologies, the software development lifecycle resides on cloud systems, and system development is entirely up to the cloud provider and not as much on the organization. Cloud computing changes the whole relationship of the systems development lifecycle to the software development lifecycle.

Shifting Left

The phrase *shift left* or *move to the left* describes the desire for security and privacy professionals to be involved as early as possible (as far to the left on a timeline) in the conceptualization, requirements, and design of a system, in order to help avoid situations where basic concept or design of a system may be deeply flawed from a security or privacy perspective.

The concept of shift left is easily understood but not easily implemented. Information security has, over the past several decades, earned the reputation of being an obstacle

to the business, so the desire to be involved earlier in the design or conceptual stage of a project requires a high level of trust on the part of business and product owners. But when cybersecurity has a "seat at the table" in earlier phases of software and product development, security can contribute to the process by opining on the ideas and concepts being developed.

The most compelling reason for cybersecurity being involved earlier is the avoidance of rework, in situations where the ideas and concepts include unsafe components.

Tying Requirements to Controls and Features

During the requirements phase of a project, organizations often group requirements into categories and assign priority levels to individual requirements. As analysis and discussion of requirements proceeds, eventually the organization arrives at a final set of requirements.

As I stated earlier, SMART requirements help organizations develop better requirements that are verifiable. Some requirements, however, may be implemented outside of the system itself, for instance, in supporting IT infrastructure. A risk assessment or threat model on the new system's design should underscore the need for security safeguards both within and outside of the system. For instance, a requirement stipulating MFA authentication will involve IT personnel in identity and access management to build their side of the required integration with a central identity service, such as LDAP or AD, in the form of reduced sign-on or single sign-on.

As described in Chapter 3, controls are usually enacted as a result of a risk assessment or incident. But it's also possible that a control may be enacted in response to one or more requirements in a system or software development project.

SDLC Risks

Most of the risk factors inherent to the SDLC affect both the entire model and each of the different phases separately. Since the SDLC could be considered to have several different facets, including a management process, an engineering framework, and even a methodology, you could say that risk factors that affect the SDLC are also common to other types of processes. For instance, earlier, I discussed risk factors that affect project and program management. These risk factors affect scope, schedule, and cost. The SDLC suffers from similar risk factors because each of the phases, as well as the SDLC, has different scope, schedule, and cost elements. Each of these phases can also be affected by risk factors that impact the availability of resources (manpower, funding, equipment, facilities, and so on), such as management commitment and organizational structure.

Risk Factors

Unique risk factors that affect the SDLC include those that affect system changes, system configuration throughout the lifecycle, and those that affect releasing different versions of the system as updates take place. Most of these elements occur during implementation, maintenance, and sustainment of a system, but they can also occur in other phases of the

SDLC where multiple versions of a system are fielded simultaneously, as well as in models where there is a continuous iterative process of updating and upgrading a system or software.

Risk factors that affect system changes include how well the change and configuration management processes work and how changes are implemented into the system. Change and configuration management processes that are not well managed are risk factors, because this could cause changes that are not documented or not tested for security, functionality, or performance. Additionally, factors such as the complexity of the system, how the organization is structured to manage the system development process, and how rapid system changes occur are also potential risk factors.

Threats

Threats to the SDLC manifest themselves when systems are produced that are not interoperable or compatible with other systems and their environment, or when systems are not secure. Other threats might include systems that do not meet the requirements they were originally intended to meet. These threats would affect both the functionality and performance of the system as well. Threats to the design process include faulty design specifications as well as faulty requirements input from the previous requirements phase. Most of the threats to the SDLC come in the form of mismanagement, miscommunication, incorrect expectations, and a lack of commitment of resources or expertise.

Vulnerabilities

Many different vulnerabilities can affect the SDLC; again, a great many of them have been previously discussed as also affecting the project or program management process. These vulnerabilities affect the scope, schedule, and cost, of course, but can also affect the quality of the system or product. In addition, the system's security, functionality, and performance are also affected by these different vulnerabilities.

While probably too many to list here, examples of vulnerabilities that affect the SDLC start with a lack of firm requirements and mixed expectations from the different system stakeholders. A lack of documentation in all phases is also a significant vulnerability to the entire SDLC, but particularly to the requirements and design phases. Failure to develop solid design specifications and systems architecture is a vulnerability that can result in a faulty design that does not meet system requirements.

In the development phase, a lack of adequate consideration for system interfaces, as well as how subcomponents fail to support higher-level components or meet technical requirements, can result in a shoddy system or software. Failure to test systems properly and thoroughly, as well as their interaction with their environment, may mean that serious functional or performance issues with the system are not discovered until well after it is put into production.

Vulnerabilities in the implementation phase include faulty implementation, lack of traceability to original requirements, and failure to properly maintain a sustained system after it has been put into production. And lastly, vulnerabilities associated with the disposal or retirement phase of the SDLC can result in systems that are not properly replaced or not securely disposed of. This can cost the organization money and potential liability.

Secure Development Practices

Secure development practices ensure that security controls, validation activities, and risk management considerations are embedded throughout the lifecycle of a system. These practices help prevent vulnerabilities, reduce defect propagation, improve code quality, and maintain alignment with organizational security requirements. Secure development integrates policy, automation, testing, and continuous assurance mechanisms rather than treating security as a discrete final step, like a final coat of paint on a house or business.

Effective secure development practices reduce risk across the lifecycle by preventing vulnerabilities early through better requirements and design, enabling consistent control implementation, and ensuring alignment with architectural, operational, and regulatory requirements.

Security Requirements and Design Controls

Security must be defined at the requirements phase and reflected in system design. Activities include specifying authentication and authorization models, data protection needs, logging requirements, resilience expectations, and compliance obligations. Architectural designs should incorporate least privilege, separation of duties, secure defaults, threat modeling results, and mitigation strategies for identified design risks. Early-stage validation reduces costly redesign and prevents fundamental architectural weaknesses.

Organizations undertaking many system development projects often develop a library of requirements to be imposed on all new projects. These requirements span security, privacy, technical standards, and other areas that do not change frequently and are common to most or all systems.

Threat Modeling

Threat modeling identifies potential attack vectors, data flows, trust boundaries, and misuse scenarios before coding begins. This enables informed design decisions, prioritization of controls, and focused testing. Common methodologies include STRIDE, attack trees, and asset–threat–vulnerability analysis. Regular threat modeling prevents security blind spots as requirements evolve.

Code Quality and Secure Coding Standards

Secure coding standards define practices that prevent common vulnerabilities such as injection, insecure deserialization, buffer errors, and improper error handling. Standards may be based on OWASP (at `www.owasp.org`), SEI (at `sci.cmu.edu/secure-development/`), or language-specific guidelines. Reviews ensure input validation, output encoding, secure API usage, safe cryptography, and proper session handling. Consistent adherence reduces vulnerability density and improves maintainability. Several code analysis tools are available for developers that integrate into their development tool suites.

Static Application Security Testing

Static application security testing (SAST) tools analyze source code or compiled binaries to identify insecure patterns, coding flaws, and logic errors. SAST is integrated early, and scans run frequently to detect issues before code merges and overall system testing. Effective programs establish rule sets aligned with organizational policies and prioritize remediation based on severity and exploitability. SAST complements DAST by examining issues with source code not visible during runtime.

Dynamic Application Security Testing

Dynamic application security testing (DAST) evaluates running applications to detect vulnerabilities such as authentication weaknesses, insecure configurations, and input handling flaws. It simulates attacker behavior without requiring source code. DAST validates runtime behavior, error messages, security headers, and exposed interfaces. It complements SAST by identifying issues that only appear during execution.

Software Composition Analysis

Software Composition Analysis (SCA) tools identify vulnerabilities in third-party libraries, open-source components, and package dependencies. They detect known CVEs, licensing conflicts, and outdated versions. Because dependency risk is a major source of exploitable weaknesses, SCA is essential for maintaining secure supply chains and ensuring that components meet organizational governance requirements.



I refer you to a seminal whitepaper by Chinmayi Sharma entitled "Tragedy of the Digital Commons," which explodes the myth that open-source software is secure because so many people are examining the code. It turns out that almost no one is. The paper is available at `https://ssrn.com/abstract=4245266`.

Code Review and Peer Validation

Manual code review supplements automated testing by evaluating logic flows, architectural consistency, and compliance with secure design principles. Reviewers examine threat modeling outputs, assess security-critical modules, and ensure that fixes do not introduce new defects. Peer review strengthens accountability and reinforces secure coding practices.

Container and Artifact Security

For containerized environments, development teams must secure base images, apply minimal OS layers, enforce scanning policies, and ensure the integrity of registry artifacts. Controls include signed images, vulnerability scanning, and continuous monitoring of image drift. Artifact integrity directly affects the security of CI/CD pipelines and deployment environments.

Secrets and Credential Management

Secure development requires proper handling of secrets, such as API keys, passwords, tokens, and certificates. Secrets must not be stored in code repositories or configuration files. Instead, teams use dedicated secret-management tools, rotation policies, and controlled access. Poor secret hygiene remains a leading source of compromise in modern development environments.

Logging, Monitoring, and Telemetry Requirements

Applications and their supporting infrastructure must implement structured logging, audit trails, and meaningful security events. Development teams define what must be logged, where logs are stored, how long data is retained, and how logs integrate with monitoring tools. Clear telemetry enables rapid incident detection and effective forensic analysis. Log data should be sent in real time to the organization's security information and event management system (SIEM) for rapid issue escalation and remediation.

Continuous Improvement and Governance

Secure development requires governance mechanisms that define standards, enforce policies, measure compliance, and require periodic review. Metrics such as vulnerability density, remediation timelines, and test coverage support continuous improvement. The concept of continuous improvement is so important that it is a required characteristic of any organization aspiring to certify its cybersecurity program to ISO/IEC 27001.

Governance ensures visibility and control on the part of senior management, as well as coordinated participation among development, operations, and security teams.

Emerging Technologies

New or emerging technologies can present risks to organizations simply because there are several important considerations when integrating the latest technologies into the existing infrastructure. Many organizations have an unfortunate tendency to rush out, buy, and attempt to implement the latest and greatest technologies on the market without careful planning or consideration of the existing infrastructure and how it will react to the new technologies. One of the primary considerations that organizations must look at is making a business case for the new technology, which may be to fill a gap that the older technology does not provide or to provide a capability that the organization must now have to compete in the market space. In other words, the organization really must justify new technologies to make them worth the risk they could incur.

Risk Factors

Emerging technologies have several risk factors inherent to their integration into the existing infrastructure. If the organization has been able to justify the implementation based on a true business need for the emerging technology, it must consider several risk

factors. One of the major risk factors is interoperability with existing infrastructure. Frequently, newer technologies don't always work properly with older systems right out of the box; adjustments may need to be made to the existing infrastructure to integrate the new technology, or even bridging technologies may be needed to connect the two together. Interoperability doesn't just involve the right connections; it can involve data formats and flows, security methods, interfaces into other systems, and changes to business processes. While regulations address uses of data (especially PII) and technology, tech regulations often lag the introduction of new technology by years. These considerations, and many others, are risk factors that must be considered before acquiring and integrating new technologies.

Security

Another risk factor is security. New technologies may have security mechanisms that are not necessarily backward compatible with existing ones. Examples include encryption algorithms and strengths, identification and authentication technologies, integrity mechanisms, and even redundant or backup systems. Additionally, the systems may involve a learning curve that may intimidate users who must now learn new security methods for the system. The human factor can be a weak link in the security chain, so either a lack of training or a lack of adaptability to the new security mechanisms can introduce risk.

System Maintenance

Earlier, when I discussed the SDLC model, I pointed out that system updates and changes can be risky if not managed properly. Integrating new technologies into the environment with older ones can introduce both intentional and unintended changes into the environment as well, affecting the stability of the organization's SDLC with a particular system. Therefore, change is also a risk factor. Even in the disposal or replacement phase of the SDLC, introducing new technologies to replace older ones can be problematic if not planned and executed properly. New technologies that are not adequately tested for functionality, performance, integration, interoperability, and security may not be able to adequately replace older systems, resulting in extended costs and possibly even requiring an extension of the older systems' lifecycles.

Threats

Threats resulting from the introduction of new and emerging technologies into the existing infrastructure are numerous. If the organization has failed to adequately plan for the new technology, these threats can become significant. Some of these threats include untested or unproven technologies, non-interoperability, incompatible security mechanisms, and the suitability of the technology for use in the organization. The organization could also incur additional costs and require more resources due to faulty implementation. These threats can be minimized through careful planning and by integrating new technologies using a stable SDLC model.

Vulnerabilities

As with threats, vulnerabilities associated with emerging technologies are numerous as well. Vulnerabilities could include a lack of trained staff committed to managing and implementing the new technology. Lack of adequate project planning is also a serious vulnerability that could affect the organization's ability to effectively integrate new technologies. Another vulnerability could be a weak support contract or another type of warranty, guarantee, or support for a new technology or system. Most of these vulnerabilities appear when a new technology is first implemented and tend to become mitigated or lessened as the technology is integrated into the existing infrastructure, but they still exist.



The key areas of concern with emerging technologies are interoperability and compatibility. These two concerns affect the security, functionality, and performance of new technologies when they are installed into an existing infrastructure.



AI is the elephant in the emerging technology room. The introduction of AI is transformational, both in terms of business improvements and in terms of new types of impact on organizations using it. AI governance is too large a topic to tackle here; refer to this author's book on AI governance (*AIGP Study Guide*, Wiley Publishing) for comprehensive coverage of this topic.

Cloud and Virtualization

Cloud and virtualization technologies abstract computing resources to improve scalability, flexibility, and efficiency. These models shift portions of operational responsibility to service providers and shift risk considerations related to configuration management, identity, data protection, and dependency on external platforms. Effective governance requires understanding how virtualized and cloud-hosted resources are provisioned, secured, monitored, and integrated into enterprise architecture.

Cloud and virtualization technologies provide powerful capabilities but require disciplined governance, identity control, secure configuration, and continuous monitoring to ensure that flexibility and scale do not introduce unmanaged risk.



The Cloud Does Not Perform All Cybersecurity Functions

With SaaS, PaaS, and IaaS cloud services being around for more than two decades, customer organizations are *still* in a state of confusion or misunderstanding about cybersecurity. Many things can and do go wrong when organizations move their IT systems to the cloud.

Virtualization Concepts and Risks

Virtualization enables multiple systems to run on shared physical hardware through hypervisors and virtual machines. Key risks include:

- Hypervisor compromise leading to full system exposure.
- VM sprawl, where rapid provisioning results in unmanaged instances.
- Inconsistent baselines, causing configuration drift and vulnerability propagation.
- Resource contention, affecting availability if hosts are overcommitted.

Proper isolation, hardened hypervisors and guests, template control, and lifecycle management reduce these risks.

Cloud Service Models (IaaS, PaaS, SaaS, XaaS)

Cloud services are delivered across layered models:

- ***Infrastructure-as-a-Service (IaaS)*:** Provides virtualized compute, storage, and networks. Customers manage OS, configurations, and workloads.
- ***Platform-as-a-Service (PaaS)*:** Provides managed runtime and development tools. Customers manage applications; providers manage infrastructure.
- ***Software-as-a-Service (SaaS)*:** Provides complete applications delivered over the Internet, with minimal customer administrative control.
- ***Everything-as-a-Service (XaaS)*:** Includes specialized offerings such as FaaS, DaaS, and managed security services.

Risks vary by model and include loss of visibility, insufficient logging, misconfigurations, data residency constraints, and provider outages.

Shared Responsibility Model

The shared responsibility model defines which security and operational tasks fall to the provider and which remain with the customer. Typically:

- **Providers** secure the underlying infrastructure, physical facilities, and core services.
- **Customers** secure identities, data, workloads, configurations, and network controls within their tenancy.

Misunderstanding shared responsibilities is a major cause of cloud-related incidents, especially misconfigured storage, open APIs, and inadequate access controls.

Shared responsibility models are depicted visually in Figure 3.1 in Chapter 3.



The details of each cloud service provider's shared responsibility models vary from one another—they are not identical. Organizations must understand these models in detail and implement controls as warranted.

Identity and Access in Cloud Environments

Identity becomes the primary security perimeter in cloud environments. Risks include:

- Excessive privilege grants
- Inadequate MFA enforcement
- Misconfigured role-based access controls
- Stale service accounts and tokens
- Overly permissive API keys or access policies
- Unprotected databases and datastores

Centralized identity governance, least privilege, continuous access review, and privileged access management reduce exposure.

Virtual Networks and Segmentation

Cloud virtual networks require explicit definition of routing, segmentation, firewall rules, and connectivity. Risks include improperly scoped security groups, overly broad network policies, unsecured VPNs, and misconfigured peering arrangements. Segmentation and strict inbound/outbound controls are essential for preventing lateral movement between systems.

Workload Protection and Configuration Management

Cloud workloads rely heavily on automated templates such as Infrastructure-as-Code (IaC), containers, or serverless functions. Risks include insecure or incomplete templates, outdated images, misconfigured services, and inadequate patching. Organizations must enforce configuration baselines, scan IaC artifacts, and validate deployment pipelines to prevent consistent replication of insecure states. Vulnerability management and asset management must include all cloud-based assets.

Data Protection and Storage Risk

Cloud storage introduces risks related to encryption, access policy errors, cross-region replication, and metadata exposure. Required controls include encryption in transit and at rest, versioning, immutable storage options, access logging, and backup validation.

Provider Dependency and Resilience

Cloud adoption introduces dependency on provider availability, SLAs, regional resilience, and service limits. Risks include:

- Regional outages
- Vendor lock-in
- API or service updates affecting compatibility
- Insufficient cross-region or multi-cloud failover
- Data compromise through cloud provider negligence

Resilience strategies include multi-zone deployments, automated failover, and documented recovery procedures.

Modern Architectures

Modern architectures emphasize modularity, scalability, automation, and distributed design. These patterns improve agility and resilience but introduce new risks related to service dependencies, configuration complexity, and identity-centric trust models. These new capabilities increase complexity and make it more difficult for organizations to know that their environments remain secure. Effective governance requires understanding how these architectures operate, how components interact, and where risks emerge across distributed environments.

Modern architectures improve performance and operational efficiency, but rely on strict configuration control, identity-centric security, and continuous validation to ensure resilience and reduce distributed risk across components and services.

AI Models

As this book is being written, the hype of AI that can solve all business problems (and beyond) is running at a fever pitch. Still, with effective AI governance and careful planning, AI and machine learning (ML) can bring considerable value to an organization in terms of operational efficiency and service quality.



Organizations must take care when designing AI systems, particularly regarding the data they select for AI system training. AI systems use data in novel ways, which sometimes conflict with established data management practices.

Quantum Computing

Quantum computing is not quite in the mainstream, but the prospects of both transformation and risk are becoming evident. From a cybersecurity and privacy perspective, the most significant risk lies in the promised ability for quantum computing algorithms to quickly defeat modern cryptography. Organizations must immediately begin to plan to upgrade their cryptosystems with *quantum-resistant cryptography*, which in some instances may take considerable time in system and application design changes and migrations.

Containerization

Containerization packages applications and their dependencies into isolated units running on a shared host OS. This improves portability and consistency across environments. Key risks include:

- Insecure or outdated base images
- Container escape and privilege escalation
- Unrestricted inter-container communication (isolation should be the default)
- Weak registry security or unsigned artifacts
- Drift between development and production images

Controls include image scanning, signed artifacts, hardened container runtimes, minimal images, configuration change detection, and strict network segmentation.

Kubernetes and Orchestration

Kubernetes automates container deployment, scaling, and management. Its complexity introduces operational and security risks, including:

- Misconfigured cluster roles, service accounts, or RBAC policies
- Open dashboards or insecure API servers
- Pod-to-pod lateral movement due to overly permissive network policies
- Inadequate logging and auditing
- Inadequate workload isolation
- Risks from admission controllers and unverified manifests

Governance requires defined RBAC, network policies, pod security standards, pod and system hardening, admission control enforcement, and monitored cluster configurations.

Microservices Architectures

Microservices divide applications into small, independently deployable services communicating through APIs. This improves scalability and development velocity but increases architectural complexity. Risks include:

- Cascading failures across dependent services
- Distributed transaction inconsistencies
- Versioning and compatibility issues
- Expanded attack surface across many endpoints
- Monitoring gaps due to decentralized logic

Resilience patterns such as circuit breakers, service mesh policies, and consistent API governance mitigate these risks.

API-first Architectures and API Security

API-first design treats APIs as primary integration points for systems and services. Strong API governance is necessary due to expanded exposure.

Risks include:

- Broken authentication and authorization
- Excessive data exposure
- Unrestricted resource consumption (rate-limiting failures)
- Insecure API gateways or misconfigured proxies
- Poor version control leading to orphaned or unmaintained endpoints

Controls include OAuth/OIDC, schema validation, rate limiting, input sanitization, API gateways, and centralized API inventories.

Zero-trust Architecture (ZTA)

Zero-trust enforces continuous verification, least privilege, and explicit access authorization regardless of network location. It supplements perimeter-based trust models.

Risks addressed by ZTA include:

- Lateral movement
- Unauthorized devices on the network
- Compromised credentials
- Overly trusted internal networks

Risks introduced by ZTA include:

- Dependency on identity systems
- Misconfigurations in policy engines
- Increased operational complexity

ZTA requires strong identity governance, continuous monitoring, granular segmentation, and validated policy enforcement.

Edge Computing

Edge computing is a distributed computing model in which computation and data storage are brought closer to end users to reduce latency. Edge computing brings certain advantages, but also introduces new risks. Unless an organization performs a risk assessment and threat modeling on proposed edge computing architectures, organizations may be altering the security posture of their environments in unexpected ways.

IT/OT Convergence

To improve efficiency, many organizations are consolidating their IT and OT environments into a single, larger environment. Operational technology (OT) environments, such as public utilities' monitoring and control of power generation and transmission infrastructure, or manufacturing automation, are notoriously less secure than IT environments, partly because they often contain unsupported components that are more vulnerable to attack. Even before the convergence of IT and OT, the security of many OT environments is already inadequate; convergence with IT may make matters far worse.

Serverless and Function-based Computing

Serverless architectures run code in managed, event-driven environments without infrastructure management. Benefits include scalability and reduced operational overhead.

Risks include:

- Overprivileged function roles
- Uncontrolled event triggers

- Short-lived, difficult-to-monitor execution environments
- Vendor lock-in and platform dependency

Policies must define least-privilege roles, event validation, secure build pipelines, and consistent logging.

Summary

Enterprise architecture within an organization affects the business's risk in several ways. Aspects of enterprise architecture risk include interoperability, supportability, security, maintenance, and how the components of the infrastructure fit within the systems development lifecycle.

Key areas in IT operations management include server and infrastructure management, end-user support, help desk, problem escalation management, and, of course, cybersecurity. IT service management models, such as ITIL and ISO/IEC 20000, guide organizations seeking to develop standard process models.

Organizations use project, program, and portfolio management to oversee and sustain both short- and long-term aspects of systems and processes.

At the core of a project are three primary drivers. These are *scope* (the amount or range of work), *schedule* (when the work is to be done and the project completion date), and *cost* (including all resources expended to complete the project).

IT service continuity management and *disaster recovery planning* are primarily concerned with the rapid response required to preserve lives, equipment, and data immediately following a serious adverse event.

Resilience refers to a business's ability to withstand disruptive events and continue its mission and operations.

Data governance and *data lifecycle management* are concerned with the creation, protection, use, retention, and disposal of business information.

The *systems development lifecycle*, or SDLC, is essentially the entire life of a system or product, from concept to development, from implementation to disposal. Security should be integrated into every step of the SDLC to ensure *security by design*.

New or emerging technologies can pose risks to organizations because several important considerations arise when integrating them into an existing environment.

Exam Essentials

Understand how enterprise architecture (EA) influences IT risk. EA defines how business processes, applications, data, and technology layers interoperate. Misalignment creates interoperability issues, technical debt, inconsistent controls, and increased operational risk. Understand the roles of reference architectures, architectural governance, and frameworks such as TOGAF and the Zachman Framework.

Recognize the risk implications of IT infrastructure components. Platforms, operating systems, databases, networks, and cloud services each introduce interoperability, security, supportability, and lifecycle risks. Heterogeneous environments increase complexity; homogeneous environments increase shared-vulnerability exposure.

Know the core IT operations management functions and their risk drivers. Service desk, incident management, problem management, change management, configuration management, monitoring, and asset management all influence service stability and control effectiveness. Weak processes in any area increase the likelihood of outages, configuration drift, or delayed failure detection.

Differentiate among project, program, and portfolio management. Projects have defined scope, schedule, and cost; programs are long-term collections of related projects; portfolios represent enterprise oversight across programs. Understand how risk factors affect each dimension and how scope creep, inadequate requirements, and weak governance create vulnerabilities.

Understand Disaster Recovery Management (DRM) and its relationship to business continuity. DRM restores systems and data after a disruption, while business continuity ensures end-to-end business operations continue. Know how RTO, RPO, RCapO, and RCO define recovery expectations. DR strategies include backup-based recovery, replication, redundant architectures, alternate sites, cloud-based failover, and rebuild-from-code approaches.

Know the types and objectives of DR testing. Readthrough, walkthrough/tabletop, communication tests, documentation checks, simulation, parallel, and full interruption tests provide increasing rigor. Testing validates the accuracy of documentation, recovery capability, dependency identification, and personnel readiness.

Understand the data lifecycle and associated governance controls. Data classification, ownership, retention, lineage, provenance, and acceptable use policies define how data must be handled. Identify the lifecycle phases: creation, storage, access, transmission, processing, sharing, retention, archiving, and deletion. Weak governance increases risks to confidentiality, integrity, availability, and compliance.

Understand the phases of the systems development lifecycle (SDLC) and the inherent risks associated with each phase. Initiation, requirements, design, development, testing, implementation, and disposal each introduce specific risk factors. Inadequate requirements, poor design, insufficient testing, weak change management, and improper retirement practices create systemic vulnerabilities.

Recognize secure development practices and their role in reducing risk. Shift left, threat modeling, secure coding standards, SAST, DAST, SCA, peer review, secrets management, and telemetry requirements improve software quality and reduce defect propagation. Governance ensures consistency and continuous improvement.

Know how modern delivery models alter risk. Agile, Scrum, DevOps, DevSecOps, CI/CD, and Infrastructure-as-Code increase speed and automation but require strong control integration. Risks include incomplete documentation, insecure pipelines, misconfigured templates, and automated defect propagation.

Identify key risks associated with emerging technologies. New technologies introduce interoperability, security, supportability, and maintenance risks. Understand how inadequate testing, lack of skillsets, weak vendor support, or incompatible security mechanisms increase exposure.

Understand how cloud and virtualization affect risk management responsibilities. Virtualization introduces risks of hypervisor, VM sprawl, and configuration drift. Cloud services rely on shared responsibility models, where misconfigurations, identity weaknesses, and provider dependencies can lead to outages or breaches.

Review Questions

The following questions are designed to test your understanding of this chapter's material. For more information on how to obtain additional questions, see this book's Introduction. Answers can be found at the end of this chapter.

1. You are managing a project that involves the installation of a new set of systems for the accounting division of your organization. You have just been told that there have been budget cuts, and the project team will not be able to purchase additional equipment needed for the installation. You must now identify additional areas to cut to fund the extra equipment. Which element of project management is most affected by this threat?
   - **A.** Schedule
   - **B.** Reporting
   - **C.** Scope
   - **D.** Quality

2. An organization is integrating risk management practices into its systems development lifecycle (SDLC). During which SDLC phase is it MOST critical to formally validate that risk treatment options are aligned with business objectives and risk appetite before significant resources are committed?
   - **A.** Requirements definition, when business needs are documented
   - **B.** System design, when architectural and control decisions are finalized
   - **C.** Development, when code and configurations are implemented
   - **D.** Testing, when control effectiveness is verified

3. During a risk assessment, management discovers that sensitive data is being copied into test environments and retained after testing is complete. Which data lifecycle control would MOST effectively address this risk?
   - **A.** Strengthening access controls on production systems
   - **B.** Increasing backup frequency for test environment data
   - **C.** Implementing data minimization and timely disposal in non-production environments
   - **D.** Encrypting all data stored in production databases

4. Which of the following is the major risk factor associated with integrating new or emerging technologies into an existing IT infrastructure?
   - **A.** Security mechanisms
   - **B.** Data format
   - **C.** Vendor supportability
   - **D.** Interoperability

5. Which of the following is a short-term process primarily concerned with protecting personnel, facilities, and equipment immediately following a disaster or major incident?
   - **A.** Disaster recovery
   - **B.** Business continuity
   - **C.** Business impact analysis
   - **D.** Recovery point objective

6. Which of the following best explains why enterprise architecture (EA) reduces operational and technology risk?
   - **A.** It documents firewall configurations for all network segments
   - **B.** It standardizes technology layers and aligns them with business processes
   - **C.** It enforces uniform programming languages across the enterprise
   - **D.** It eliminates the need for system lifecycle management

7. As part of data lifecycle management, an organization is reviewing how data is handled after decommissioning a core business application. Which action is MOST important to reduce long-term risk while maintaining regulatory compliance?
   - **A.** Retain all application data indefinitely in primary storage
   - **B.** Immediately destroy all associated data to reduce storage costs
   - **C.** Transfer all data ownership to the IT operations team
   - **D.** Archive data according to defined retention and disposal requirements

8. A heterogeneous platform environment increases risk primarily because it:
   - **A.** Requires a single vendor for support contracts
   - **B.** Reduces the number of required security controls
   - **C.** Increases complexity and interoperability challenges
   - **D.** Eliminates the need for patch management processes

9. Which IT operations process most directly reduces the likelihood of configuration drift across systems?
   - **A.** Incident management
   - **B.** Configuration management
   - **C.** Portfolio management
   - **D.** Problem management

10. A project experiences repeated delays due to unclear definitions of required deliverables. Which project element is most directly at risk?

A. Cost

B. Quality

C. Scope

D. Staffing

11. Which disaster recovery metric defines the maximum tolerable amount of data loss measured in time?

A. RTO

B. RCapO

C. RPO

D. MTTF

12. An organization is introducing an IoT-enabled manufacturing system. The vendor requires multiple external connections to cloud APIs. During the risk assessment, the IT team identifies that the new system uses a proprietary authentication method incompatible with the enterprise identity service. What is the MOST significant risk associated with integrating this emerging technology?

A. Increased bandwidth consumption across the manufacturing network

B. Failure of the system to meet RCapO requirements

C. Interoperability and security integration gaps with existing infrastructure

D. Inability to migrate historical data during system upgrades

13. Which IT operations function addresses systemic causes of recurring incidents?

A. Incident management

B. Problem management

C. Vendor management

D. Lifecycle management

14. Which emerging technology risk is most commonly associated with integrating new systems into older enterprise environments?

A. Increased storage efficiency

B. Simplified authentication processes

C. Interoperability and compatibility failures

D. Guaranteed backward-compatible security controls

15. A large enterprise experiences recurring incidents where unauthorized configuration changes cause outages in its virtualized server environment. Investigation reveals that VM templates are being manually altered by several administrators without documentation. Which control would MOST effectively reduce this risk?

A. Increasing the frequency of DR tests

B. Using immutable, version-controlled templates managed through configuration management

C. Adding additional virtualization hosts to distribute workloads

D. Implementing annual cybersecurity awareness training

16. Which data governance activity ensures that only the correct individuals continue to have access to sensitive datasets?

A. Data provenance

B. Access review

C. Data lineage analysis

D. Database indexing

17. Which type of DR test provides the highest assurance but also the greatest operational risk?

A. Full interruption test

B. Parallel test

C. Readthrough test

D. Tabletop test

18. Which of the following creates the highest risk in a CI/CD pipeline?

A. Frequent code deployments

B. Storing secrets in build scripts

C. Using multiple automated test suites

D. Requiring peer reviews before merging code

19. Which practice best represents a "shift left" approach in a secure development process?

A. Conducting penetration tests annually

B. Executing DR tests in a separate environment

C. Performing threat modeling before design is finalized

D. Implementing MFA during system retirement

20. Alternate processing sites, such as hot or warm sites, are primarily used to satisfy which requirement?

A. Minimizing the cost of vendor contracts

B. Meeting RTO and maintaining service availability

C. Reducing the number of DR tests required

D. Eliminating the need for data backups

Answers to Review Questions

1. C. Scope is the element of a project most affected by insufficient funding. The project will need to shrink its scope to stay within budget.

2. B. The system design phase is the point at which architectural choices, control selection, and risk treatment decisions are concretely defined. Validating alignment with business objectives and risk appetite at this stage prevents costly rework and ensures significant risks are addressed appropriately before committing development resources. Requirements definition identifies needs but does not finalize treatments, while development and testing occur after key risk decisions have already been made.

3. C. Applying data lifecycle controls, such as data minimization and timely disposal, ensures that sensitive data is not retained longer than necessary, particularly in non-production environments. This directly reduces exposure risk. The other options do not address unnecessary retention outside the intended lifecycle stage.

4. D. Interoperability is the primary risk when integrating new or emerging technologies into an existing IT infrastructure. Interoperability is the best answer, as it encompasses the individual factors in the other answers.

5. A. Disaster recovery is a short-term process primarily focused on protecting personnel, facilities, and equipment immediately after a disaster or major incident.

6. B. EA reduces risk by creating a standardized, governed structure linking business processes, data, applications, and technology. This reduces interoperability failures, technical debt, and inconsistent control implementation.

7. D. Data lifecycle management requires that data be retained, archived, or disposed of in accordance with legal, regulatory, and business requirements. Archiving in accordance with defined retention and disposal requirements balances compliance obligations with risk reduction. Indefinite retention increases risk, while immediate destruction may violate regulations.

8. C. Supporting multiple operating systems and architectures increases interoperability challenges, support overhead, and configuration complexity, all of which directly contribute to operational and security risks.

9. B. Configuration management maintains authoritative configuration records, baselines, and relationships, allowing consistent control of system states and preventing (or detecting) drift.

10. C. Unclear or poorly defined deliverables indicate an incomplete scope, leading to misaligned expectations, rework, and schedule instability.

11. C. RPO (Recovery Point Objective) specifies the maximum acceptable data loss. It determines the required backup frequency or replication strategy.

12. C. Proprietary authentication that is incompatible with enterprise identity systems creates major interoperability and security risks, including weakened access control, inconsistent policy enforcement, and an increased attack surface.

13. B. Problem management conducts root cause analysis (RCA) to identify and address the underlying causes of recurring incidents and reduce operational risk.

14. C. Emerging technologies often introduce challenges involving protocols, data formats, authentication, and interfaces.

15. B. Uncontrolled changes to VM templates cause configuration drift and unplanned outages. Enforcing template version control, baseline integrity, and configuration management directly mitigates this risk.

16. B. Access reviews confirm that users, groups, and roles still require the access they hold, reducing the risk of unauthorized access.

17. A. A full interruption test shuts down production systems and performs a full recovery, providing maximum assurance, but it also introduces substantial operational risk if the recovery system fails.

18. B. Hard-coded secrets in pipelines expose credentials and other sensitive information, enabling compromise of build systems, repositories, or production environments.

19. C. Shift left moves security earlier in the lifecycle: threat modeling before design is finalized ensures risks are addressed before development.

20. B. Alternate sites ensure critical systems can be restored quickly, directly supporting required RTOs and service resilience.



Chapter 5

Information Security

This chapter covers CRISC Domain 4, "Information Technology and Security," specifically the Information Security Principles subdomain. The job practice for this domain is described as follows:

In this domain, we interrogate the alignment of business practices with Risk Management and Information Security frameworks and standards, as well as the development of a risk-aware culture and implementation of security awareness training.

The preceding statement from ISACA describes the whole of Domain 4. This chapter covers the information security portion of this domain, while Chapter 4 covers the information technology portion.

This chapter covers these job practice elements:

✔ B—INFORMATION SECURITY PRINCIPLES

1. *Information Security Concepts, Frameworks, and Standards*
2. *Information Security Awareness Training*
3. *Business Continuity Management*
4. *Data Privacy and Data Protection Principles*

A single CRISC Task Statement relevant to this domain focuses on risk treatment and reporting. It is:

8. ***Promote a risk-aware culture by contributing to the development and implementation of security awareness training.***

The topics in Chapter 4 and this chapter represent 20% of the CRISC examination.

Information security plays a central role in managing risk to modern organizations. As enterprises increasingly depend on information and technology to achieve strategic and operational objectives, failures to adequately protect information assets can result in regulatory penalties, financial loss, operational disruption, and long-term reputational harm. This chapter focuses on the information security principles within CRISC Domain 4 and examines how security practices support effective risk management and governance.

Unlike purely technical discussions of security controls, this chapter emphasizes information security from a risk-based and business-aligned perspective. It examines how security principles, frameworks, policies, and practices are used to reduce risk to acceptable levels, support organizational resilience, and enable informed decision-making by management and risk owners. The chapter also highlights the importance of aligning information security activities with risk appetite, risk tolerance, and enterprise risk management (ERM) processes.

Key topics include foundational security concepts such as confidentiality, integrity, and availability; the role of widely adopted frameworks and standards; access management and incident response practices; security and risk awareness training; business continuity considerations; and core data privacy and data protection principles and regulations. Throughout the chapter, security concepts are framed as contributors to risk identification, risk response, and ongoing risk monitoring.

For CRISC candidates, this chapter reinforces the idea that information security is not an isolated technical function, but a critical component of organizational risk management and governance.

Information Security Concepts, Frameworks, and Standards

To successfully pass the CRISC exam, you should be familiar with basic information security concepts. You can't be expected to know how to manage risk in a security environment if you don't understand the basics of information security in the first place. I'm going to assume you have some level of experience already as an information security professional, since risk management is a significant portion of (and a logical career progression from) the information security profession. You may also have had some level of experience in specific risk management processes at some point during your career. As such, I don't go into great detail on the basic information security concepts in the upcoming sections; this brief review will serve as a quick refresher to remind you of key security concepts. I also review a few key control frameworks in detail, including the NIST Risk Management Framework (RMF) and COBIT 2019, and explain how these frameworks are used to manage risk. Additionally, this chapter covers more material than you might find on the exam; this is for two important reasons. First, it's important that you understand foundational and conceptual knowledge about controls and some of the relevant control frameworks you will find on both the exam and in the real world as a risk and control assessor. Second, this book is intended not only as a study guide for the exam but also as a comprehensive desk reference you will use even after you have passed the CRISC exam.

You may learn traditional information security doctrine, as well as fundamental information security knowledge, from various training courses and on-the-job experience over the years. One such doctrine teaches that there are three fundamental information security goals. These goals are what we're striving for as risk and information security

professionals: confidentiality, integrity, and availability. You'll sometimes see these three terms strung together as an acronym, such as the *CIA triad* or, occasionally, the *AIC triad*, depending on the different information security literature you read. In any event, these three goals are what you want to achieve for all your information systems and data. They are also characteristics you want to ensure all your systems, processes, procedures, methods, and technologies display.

Managing security and risk is not an ad hoc process. It is a complex effort and involves establishing a formal program with responsible people leading it. It requires developing procedures and processes that are defined, repeatable, and defendable. Fortunately, you don't have to reinvent the wheel; most of this work has already been done for you in the form of established frameworks, methods, standards, and practices. One of the first things you'll want to do when establishing a cyber risk program is to understand what type of framework, processes, standards, and practices you will use, since there is a variety to choose from. You must try to use the one that best fits your organization, and you can't do that unless you have at least a basic understanding of the more defined, standardized ones used in the industry. Let's take a moment and discuss the difference among frameworks, standards, and practices.

Confidentiality, Integrity, and Availability

Confidentiality is the goal of keeping information systems and data from being accessed by people who do not have the authorization, need-to-know, or security clearance to access that information. In other words, confidentiality means that only authorized individuals and entities should be able to access information and systems. Confidentiality can be achieved through several security protection mechanisms, such as rights, privileges, permissions, encryption, authentication, and other access controls. If confidentiality of data or information systems is breached, we get the opposite of confidentiality, which is unauthorized disclosure. *Unauthorized disclosure* is a risk to data and information systems, and one we as security professionals struggle very hard to protect against.

Integrity is the characteristic of data that describes the state in which a system or its data has not been subject to unauthorized modification. Data integrity means that data remains in the same state as when it was stored or transmitted. Therefore, when it is accessed again or received, it should be identical to the data that was originally stored or transmitted. Integrity is achieved in several ways, by using checksums, message digests, and other verification methods. Data alteration is the opposite of integrity, particularly when the modification has not been authorized by the data owner. Data modification can happen accidentally, such as when it is inadvertently changed due to human error or faulty transmission media. It can also happen intentionally (which is usually malicious in nature when this modification is unauthorized) by direct interaction with data during storage or transmission (during an attack, for example). This risk to data affects whether the data can be trusted as authentic or genuine, whether it can be read as intended, or whether it is corrupt.

Availability is the property of systems and data to be accessible to authorized users at any time or under any circumstances. Even if data is kept confidential and its integrity is kept intact, it still does us no good if we can't access it when we need it to perform critical business functions. Availability ensures we have this data (and the information systems that process it)

at our fingertips. Just as confidentiality and integrity have their opposites, *data destruction* and *denial-of-service* attacks are the opposites of availability. These risks to our information systems could prevent authorized consumers of that data or users of that information system from performing their jobs, thus severely impacting our business operations.

Popular security theory sets forth these three overarching security goals but also provides for auxiliary elements that support these goals in various ways. These concepts, both individually and in combination, help us, as security professionals, ensure that we protect our systems and data from unauthorized use or misuse. The next few sections discuss all these different security elements and other concepts, as well as how they support three primary goals of information security.



With the advent of autonomous vehicles, robotic surgery, and more, the principle of *life safety* is an emerging topic in information security, as is the traditional CIA: confidentiality, integrity, and availability.



You will need to understand the definitions of the goals of information security—confidentiality, integrity, and availability—very well for the exam. Almost everything in information risk management supports those three goals, either directly or indirectly.

Frameworks

A *framework* is a generally overarching methodology for a set of activities or processes. It may not get into the detailed processes and procedures; instead, it provides for a 500-foot view of the general direction and steps used to build a more detailed program or process. A framework is used as an overall architecture for a greater effort. A framework has characteristics that include defined steps and repeatability, and it can be tailored based on the organization's needs. In terms of a risk management framework, you may have a set of general steps defining how to approach risk management, which include listing the processes and activities necessary to build such a program or effort. You would then break down these larger steps into specific supporting procedures for this effort based on the needs of your organization and using standards (described next). Frameworks are typically selected and adopted at the strategic level of corporate management and governance.

NIST RMF

The *NIST Risk Management Framework (RMF)* is a seven-step methodology that provides for risk management all the way through the information systems lifecycle. The steps for the RMF are briefly described here:

- **Step 1: Prepare** This step has the purpose of carrying out essential activities to help prepare all levels of the organization to manage its security and privacy risks using the RMF. It involves identifying key risk management roles, establishing the organizational risk management strategy, determining the organization's risk tolerance, assessing the

organization's risk assessment, developing the organizational strategy for continuous monitoring, and identifying the implemented common controls.

- **Step 2: Categorize** This step involves inventorying the types of information on target systems and assigning *security categorization* levels to that information based on the level of impact if the security goals of confidentiality, integrity, and availability were affected or compromised for the information on the system. This step uses subjective values of high, medium, and low to assign values to each of the three goals for a particular type of information. Types of information processed on the system could include business-sensitive, financial, protected health information, and so on. FIPS 199, "Standards for Security Categorization of Federal Information and Information Systems," as well as NIST Special Publication 800-60, "Guide for Mapping Types of Information and Information Systems to Security Categories," provide detailed guidance on categorizing information systems.
- **Step 3: Select** This step provides baselines of security controls based on the high, medium, and low values assigned during Step 2. If the aggregate value of information or system has been rated as *high*, for example, the high baseline of security controls is employed for that system. Once a security control baseline has been established, the organization has the latitude and flexibility to add or subtract security controls from the baseline as it sees fit based on different factors, including the applicability of some controls, the environment the system operates within, and so on. The selected controls can be found in the supporting NIST Special Publication 800-53, revision 5, "Security and Privacy Controls for Information Systems and Organizations," which contains a catalog of all the NIST controls.
- **Step 4: Implement** Here, the selected controls are applied to the information systems, and data is processed on those systems. This is an extensive process that can cover a good deal of the lifecycle of the system in question, and it may take significant time and resources. In this step, the organization is essentially securing the information system against any validated threats and protecting identified vulnerabilities.
- **Step 5: Assess** This step is where a lot of security professionals who manage certification and accreditation activities or perform risk assessments come into the picture. During this step, the controls the organization selects for the information system are formally assessed, verifying that they were implemented correctly and validating whether they perform as they were designed. They are assessed based on their effectiveness in protecting against the threats they were implemented to protect against. During this step, the system is assessed in its current state, with all existing controls and mitigations in place. Based on the assessment findings, there may be recommendations for further controls and mitigations, as well as alterations to the existing security posture for the system. In this step, the level of risk to the system and its data is typically analyzed and determined.
- **Step 6: Authorize** This step involves the decision from the entity in charge to authorize a system to be implemented and put into operation. This decision is based on various factors, including the level of risk assessed during Step 5, the risk appetite

the organization has settled on, and the tolerance for risk the organization is willing to accept. The decision to authorize a system for use may also come with caveats, including conditional authorization based on the continued mitigation and reduction of risk by the system or data owner. This authorization is a formal authority for the system to operate, made by someone with the legal authority to make that decision. It is typically in writing and only valid for a specified period, after which the system must be reassessed for risk and control compliance.

- **Step 7: Monitor** Continuous monitoring of security controls defines Step 7 in the RMF; just because an authorization decision is rendered doesn't mean the system will now be operated forever without continually monitoring its security posture for new or changed risks. Existing controls should be monitored for continued compliance and effectiveness against identified threats. New risks will occasionally be discovered for the system as new threats and vulnerabilities are identified, and the system will have to be reauthorized after a specific period. Note that the RMF is a cyclical process; all these steps will be accomplished again for each system at various times over its lifecycle.

NIST CSF

The *NIST Cybersecurity Framework (CSF)* is a high-level, voluntary framework developed by the NIST to help organizations manage and reduce cybersecurity risk. Released in 2024, CSF 2.0 expands the original framework beyond critical infrastructure and places explicit emphasis on governance, risk ownership, and integration with enterprise risk management (ERM). Unlike the NIST RMF, which is system-centric and prescriptive, the CSF is outcome-oriented and intended to guide cybersecurity risk management at the organizational level.

CSF 2.0 is organized around six core functions, which represent the full lifecycle of cybersecurity risk governance and management:

- **Govern:** The Govern function is new in CSF 2.0 and establishes the organizational context for managing cybersecurity risk. It includes defining roles and responsibilities, setting policies and standards, aligning cybersecurity strategy with business objectives, establishing risk appetite and tolerance, and ensuring oversight and accountability. From a CRISC perspective, this function directly supports risk governance, risk ownership, and alignment with ERM.
- **Identify:** The Identify function focuses on understanding the organization's assets, business environment, dependencies, and associated cybersecurity risks. Activities include asset management, risk assessment, and identification of threats and vulnerabilities. Identify is driven by governance decisions and provides the foundation for informed risk analysis and prioritization.
- **Protect:** Protect addresses safeguards used to manage risk and limit the potential impact of cybersecurity events. It includes access control, awareness and training, data security, information protection processes, and protective technologies. These activities represent preventive and deterrent risk responses aligned with organizational risk tolerance.
- **Detect:** The Detect function defines activities to identify the occurrence of cybersecurity events in a timely manner. Continuous monitoring, anomaly detection, and event

analysis reduce risk exposure by shortening the time between threat exploitation and organizational response.

- **Respond:** Respond focuses on actions taken once a cybersecurity incident is detected. It includes response planning, communications, analysis, mitigation, and lessons learned. From a risk management standpoint, this function supports incident-driven risk response, escalation, and decision-making.
- **Recover:** Recover addresses the restoration of capabilities or services impaired by cybersecurity incidents. It includes recovery planning, improvements, and communications. This function aligns closely with resilience, business continuity, and disaster recovery objectives.

In addition to the core functions, CSF 2.0 includes categories, subcategories, and profiles, along with mappings to other frameworks and standards such as NIST SP 800-53, ISO/IEC 27001, and COBIT. These mappings enable organizations to use the CSF as a unifying structure across multiple governance, risk, and compliance initiatives.



CRISC candidates should understand that CSF 2.0 is not a control catalog or a step-by-step methodology. Instead, it provides a common structure and vocabulary for governing and managing cybersecurity risk, enabling alignment between security activities, business objectives, and enterprise risk management practices.

ISO/IEC 27001/27002/27701/31000

Whereas the NIST RMF is very much an American framework, the International Organization for Standardization (ISO) frameworks are used globally. The ISO family is focused mainly on keeping information assets secure, ensuring privacy, and managing risk across the following primary documents:

- **ISO/IEC 27000:2018:** Provides an overview of information security management systems (ISMS).
- **ISO/IEC 27001:2022:** Specifies the requirements for establishing, implementing, maintaining, and continually improving an ISMS within the context of the organization. It also includes requirements for the assessment and treatment of information security risks tailored to the needs of the organization.
- **ISO/IEC 27002:2022:** Gives guidelines for organizational information security standards and information security management practices, including the selection, implementation, and management of controls, taking into consideration the organization's information security risk environment(s).
- **ISO/IEC 27701:2025:** Specifies requirements and provides guidance for establishing, implementing, maintaining, and continually improving a privacy information management system (PIMS).
- **ISO 31000:2018:** Provides guidelines on managing risk faced by organizations.

ISO/IEC 27001 Certification

While seeking ways of improving confidence in its information security program, a cloud hosting organization explored options and chose to pursue certification to the ISO/IEC 27001 standard. The path to certification was straightforward and consists of these steps:

- **Determine the scope of the certification.** For a service provider, this likely refers to its service operations and supporting processes. If corporate IT is structurally separate, it may be out of scope.
- **Determine the set of applicable controls.** These may be the controls found in ISO/IEC 27002, but other controls may be used.
- **Perform a gap analysis.** The organization needs to assess its current compliance with all applicable requirements and the effectiveness of its controls.
- **Undergo an audit.** The organization next solicits services from one of many accredited ISO audit firms, which will perform an audit of the organization's compliance with ISO/IEC 27001 requirements. The audit firm will also assess some, most, or all of the organization's controls.
- **Receive certification.** After the audit, the audit firm will analyze its findings and produce assessment deliverables for the cloud services provider. If the cloud services provider warrants certification, it will be awarded. It's also possible that the cloud services provider may need to make specific improvements before certification is awarded.

Maintaining ISO/IEC 27001 certification requires periodic audits (typically annual) to confirm that the customer continues to manage its security program effectively.

COBIT 2019

COBIT (formerly an acronym for Control Objectives for Information and Related Technologies) is a framework developed by ISACA that facilitates the governance of enterprise information and technology. This framework has been around for several years and through several iterations. It also provides for integration of other popular frameworks and standards, including the Open Group Architecture Framework (TOGAF), the Project Management Body of Knowledge (PMBOK), the Information Technology Infrastructure Library (ITIL), Projects IN Controlled Environments 2 (PRINCE2), the Committee of Sponsoring Organizations of the Treadway Commission (COSO), and the many International Organization for Standardization (ISO) standards. This interoperability enables new users of COBIT to leverage any of these other standards they have already been using in their adoption of COBIT.

COBIT consists of two layers in its model, governance and management, which are further broken down into a total of 40 separate objectives. Table 5.1 quickly summarizes these layers and their respective objectives.

TABLE 5.1 COBIT 2019 governance and management objectives.

Governance: Evaluate, Direct, and Monitor (EDM)				
EDM01: Ensured Governance Framework Setting and Maintenance	EDM02: Ensured Benefits Delivery	EDM03: Ensured Risk Optimization	EDM04: Ensured Resource Optimization	EDM05: Ensured Stakeholder Engagement
Management: Align, Plan, and Organize (APO)				
APO01: Managed I&T Management Framework	APO02: Managed Strategy	APO03: Managed Enterprise Architecture	APO04: Managed Information	APO05: Managed Portfolio
APO06: Managed Budget and Costs	APO07: Managed Human Resources	APO08: Managed Relationships	APO09: Managed Service Agreements	APO10: Managed Vendors
APO11: Managed Quality	APO12: Managed Risk	APO13: Managed Security	APO14: Managed Data	
Management: Build, Acquire, and Implement (BAI)				
BAI01: Managed Programs	BAI02: Managed Requirements Definition	BAI03: Managed Solutions Identification and Build	BAI04: Managed Availability and Capacity	BAI05: Managed Organizational Change
BAI06: Managed IT Changes	BAI07: Managed IT Change Acceptance and Transitioning	BAI08: Managed Knowledge	BAI09: Managed Assets	BAI10: Managed Configuration
BAI11: Managed Projects				
Management: Deliver, Service, and Support (DSS)				
DSS01: Managed Operations	DSS02: Managed Service Requests and Incidents	DSS03: Managed Problems	DSS04: Managed Continuity	DSS05: Managed Security Services
DSS06: Managed Business Process Controls				

(*Continued*)

TABLE 5.1 (Continued)

Management: Monitor, Evaluate, and Assess (MEA)			
MEA01: Managed Performance and Conformance Monitoring	MEA02: Managed System of Internal Control	MEA03: Managed Compliance with External Requirements	MEA04: Managed Assurance

Note that while COBIT covers a variety of business and IT processes and areas, those specific to risk management happen at both layers—governance and management—and are tightly integrated with other processes. COBIT, in this regard, is not a risk management framework per se, as is the NIST RMF, but offers more of a broader view of management and governance across all primary areas of a business. The next topic, ISACA's Risk IT Framework, supports COBIT and provides a more granular view of risk management practices and activities.



Understand that while COBIT is important to the overall risk discussion, it is not a risk framework itself. It does, however, support management and governance of IT as well as other critical business areas. It also leads to a more detailed discussion on the ISACA Risk IT Framework, which does deal with IT risk.

The Risk IT Framework (ISACA)

The *Risk IT Framework* is a more concise, risk-related set of processes than that offered by the related framework, COBIT 2019. While COBIT covers the "big picture" of governance and management processes that support risk management programs in the organization, the Risk IT Framework gets more into the key processes of risk management, such as risk governance, evaluation, and response. The Risk IT Framework is also a summary of ISACA's *The Risk IT Practitioner Guide*, available to ISACA members, which is an even more in-depth treatment of the processes and activities encountered during risk management.

The Risk IT Framework comes from traditional risk management principles of various ERM standards and describes activities and processes thought of as best practices in the industry. It provides a starting place for the establishment of these processes—sort of a map to navigate from a nonexistent or immature risk management program to a formalized, defined set of processes. The Risk IT Framework focuses more on the business risks of using IT in the organization's structure and how risk is involved with the gap often found between IT implementation and business goals. Table 5.2 lists the different domains of the Risk IT Framework model, with each of their three major processes.

TABLE 5.2 Domains and processes of ISACA's Risk IT Framework.

Risk Governance (RG)	RG1: Establish and maintain a common risk view	RG2: Integrate with ERM	RG3: Make risk-aware business decisions
Risk Evaluation (RE)	RE1: Collect data	RE2: Analyze risk	RE3: Maintain risk profile
Risk Response (RR)	RR1: Articulate risk	RR2: Manage risk	RR3: React to events

Keep in mind, I've only covered a few of the risk-related frameworks and standards available in the industry. Also, be aware that I've only scratched the surface of these bodies of knowledge here in this chapter; a complete discussion of each one of them is beyond the scope of this book. You should obtain and review these publications in depth before you sit for the exam. There are also many other available frameworks and standards developed and promulgated by other professional organizations, government entities, industries, vendors, and so on. You should be familiar with as many of these as are relevant to your work, since they help you turn risk management from "magic" into an art and science for your organization. The frameworks and standards I have described earlier are the ones most relevant to your studies for the CRISC exam.



While you may not be expected to know the intricate details of each framework described here, it will be helpful to know at least the basic characteristics and descriptions of each. You may be asked to identify a particular characteristic of a framework on the exam.



Frameworks and Risk Management Alignment

The frameworks and standards discussed in this chapter, including NIST RMF, NIST CSF, COBIT 2019, ISO/IEC standards, and the Risk IT Framework, approach information security from different perspectives and levels of abstraction. Despite these differences, each framework ultimately supports the same core risk management activities:

- Establishing governance and accountability
- Identifying and evaluating risk
- Selecting and implementing appropriate controls
- Monitoring effectiveness
- Responding to realized risk events

These frameworks provide structured methods for aligning information security practices with organizational risk appetite, risk tolerance, and business objectives, without prescribing a single required implementation model.

Policies

A *policy* is a management-approved statement that describes a required behavior, characteristic, or outcome. First discussed in Chapter 1, policies are often considered to resemble regulations that apply to the organization, its personnel, and its systems.

Information Security Policy

Information security policy is a corporate policy that addresses many aspects of personnel behavior, as well as the characteristics of business processes and information systems. Information security policy covers these and other topics:

- Security-related roles and responsibilities
- Risk assessments and risk management
- Security awareness training
- Internal audit
- Identity and access management
- Data management and classification
- Third-party risk management
- Human resources security
- Event logging, investigations, and incident response
- Physical security
- System, network, and device configuration and hardening
- Software and system development security
- Resilience, backup, and recovery

Organizations often organize their information security policies to resemble the structure of the security framework(s) that must be followed. Security policies (and other policies) should be formally reviewed by management at least once per year, and formally managed through a document management system to control and record reviews and updates.

Privacy Policy

A privacy policy is a corporate policy that describes the collection, use, and disposal of sensitive information about employees, customers, constituents, and others. The enactment of privacy laws, together with societal expectations, compels many organizations to develop a privacy policy, so that affected data subjects are aware of the organization's practices concerning their personal, sensitive information.

Privacy policy is discussed in detail later in this chapter.

Acceptable Use Policy (AUP)

An *acceptable use policy (AUP)* is a policy consisting of a set of established guidelines for the appropriate use of computer networks within an organization. The AUP is a written agreement, read and signed by *workers*, that outlines the organization's terms, conditions, and rules for Internet and internal network use and data protection.

An AUP helps educate workers about the kinds of tools they will use on the network and what they can expect from those tools. The policy also helps define boundaries of behavior and, more critically, specifies the consequences of violating those boundaries. The AUP also lays out the actions that management and the system administrators may take to maintain and monitor the network for unacceptable use, and it includes the general worst-case consequences or responses to specific policy violations.

Developing an AUP for your organization's computer network is extremely important for both organizational security and limiting legal liability in the event of a security issue. An AUP should cover the following topics:

- **Legality:** The organization's legal department needs to approve the policy before it's distributed for signing. The policy will be used as a legal document to ensure that the organization isn't legally liable for any type of Internet-related incident and any other transgressions, such as cracking, vandalism, and sabotage.
- **Uniqueness to your environment:** The policy should be written to cover the organization's specific network and the data it contains. Each organization has different security concerns—for example, a medical facility needs to protect data that differs significantly from that of a product sales organization.
- **Completeness:** Beyond rules of behavior, the AUP should also include a statement concerning the organization's position on personal Internet use on company time.
- **Adaptability:** Because the Internet is constantly evolving, the AUP will need to be updated as new issues arise. You can't anticipate every situation, so the AUP should address the possibility of something happening that isn't outlined.
- **Protection for employees:** If your employees follow the rules of the AUP, their exposure to questionable materials should be minimized. In addition, the AUP can protect them from dangerous Internet behavior, such as giving out their names and email addresses to crackers using social engineering techniques.

The focus of an acceptable use policy should be on the responsible use of computer networks and the protection of sensitive information. Such networks include the Internet—for example, web, email (both personal and business), social media, and instant messaging access—and the organization's intranet. An AUP should, at a minimum, contain the following components:

- A description of the strategies and goals to be supported by Internet access in the organization
- A statement explaining the availability of computer networks to workers
- A statement explaining the responsibilities of workers when they use the Internet

- A code of conduct governing behavior on the Internet
- A description of the consequences of violating the policy
- A description of what constitutes acceptable and unacceptable use of the Internet
- A description of the rights of individuals using the networks in the organization, such as user privacy
- A description of the expectations for the access and use of information
- Proper use of social media
- A disclaimer absolving the organization from responsibility under specific circumstances
- A form for workers to sign, indicating their agreement to abide by the AUP

Note that many organizations' websites contain an acceptable use policy or terms of use statement that protects them from any liability from users of the site.

Standards

A *standard* is a mandatory set of specifications, procedures, or processes used by the organization and usually fits into an overall framework. Standards often define more detailed processes or activities used to perform a specific set of tasks. Standards are used for compliance reasons and made mandatory by an organization or its governance.

The National Institute of Standards and Technology (NIST) standards are mandatory for use by the United States federal government, for instance, but are published as an option for private organizations and industries to adopt if they so choose. If an organization adopts the NIST standards for risk management, for example, the organization may make them mandatory for use by its personnel. Then all processes and activities for a given effort within the organization would have to use and meet those standards.

Some standards define the level of depth or implementation of a security control or measure. The Federal Information Processing Standards (FIPS) for cryptography and encryption are an example of this; they set forth the different levels of encryption strength for various cryptography applications that may be required in certain circumstances. So, if you create security policies and procedures for implementing cryptography within the organization, the FIPS standard could tell you to what level those policies and procedures must be implemented.



Remember that there are two types of standards: first, an internal corporate standard that specifies detailed requirements, rules, or configurations. Second, a specification, guideline, or benchmark published by a recognized standards body, such as ISO or NIST.

Controls

Controls are safeguards implemented to reduce risk by preventing, detecting, or correcting adverse events that could affect the confidentiality, integrity, or availability of information assets. In the context of information security, controls may be administrative, technical, or

physical in nature, and they are selected and designed to support business objectives while operating within the organization's risk appetite and tolerance.

From a CRISC perspective, controls are not evaluated based solely on their existence, but on their effectiveness in reducing risk. Controls are a primary means of risk treatment, but they do not eliminate risk entirely. Instead, they influence the level of residual risk that remains after mitigation. Management is responsible for determining whether residual risk is acceptable, or whether additional action is required.

It is important to recognize that control failure itself is a significant risk driver. A control may fail due to poor design, improper implementation, lack of maintenance, lack of training, unclear documentation, inadequate monitoring, or intentional or unintentional circumvention. Controls may also degrade over time as business processes, technologies, threats, or organizational structures change. In such cases, continued reliance on ineffective controls can lead to a false sense of security and an underestimation of risk exposure.

Examples of control failure include access reviews that are performed but not remediated, incident response plans that exist but are never tested, vendor security requirements that are defined contractually but not enforced, or technical controls that are misconfigured or bypassed by privileged users. In each case, the control is present, yet risk remains elevated due to ineffective operation.

Effective governance requires that control owners and risk owners be clearly identified and held accountable for monitoring control performance and addressing deficiencies. When control weaknesses are identified, management must decide whether to remediate the control, implement compensating controls, transfer the risk, or formally accept the residual risk.

A comprehensive discussion of control types, control design, control testing, and control effectiveness analysis is provided in Chapter 3, which focuses on risk response and mitigation.

Practices

A *practice* is a normalized process that has been tried and proven as generally accepted within a larger community. Practices could also be developed by a standards organization or a recognized authority regarding a particular subject or process. Professional industry organizations or vendors often develop practice documents. You might also see "best practices" published by various industries or organizations, for example. Practices are not usually mandatory, but could be made mandatory by the corporate management or other governance if they were so inclined.

The next few sections give more detailed examples of some of the formal frameworks and standards you should be familiar with for the exam and in real life as a risk management professional. I recommend you pay particular attention to the ones developed and published by ISACA; these will likely be present in some form on the exam. Of course, in this book, I give you only a brief overview of each, and you should take the time to review the actual standards and frameworks in depth before you sit for the exam.

Due Care, Due Diligence, and Due Process

Due care, *due diligence*, and *due process* are terms that apply to the implementation and enforcement of organization-wide security policies. An organization practices *due care* by taking responsibility for all activities that take place in corporate facilities. An organization practices *due diligence* by implementing and maintaining these security procedures consistently to protect the organization's facilities, assets, and workers. Although many organizations outline plans for security policies and standards, they often never officially implement them or don't properly share the information with workers. Without direction from management in the form of training, guidelines, and manuals, and without workers' input and feedback, security policies will not be successful.

By practicing due care, the organization shows it has taken the necessary steps to protect itself and its workers. By practicing due diligence, the organization ensures that these security policies are properly maintained, communicated, and implemented. If the organization doesn't follow proper due care and due diligence initiatives, it might be considered legally negligent if organizational security and customer data are compromised.

Due process guarantees that in the event of a security issue attributable to a worker, the worker receives an impartial and fair inquiry into the incident to ensure the worker's employment rights are not being violated. If during an investigation or inquiry, the worker's rights are violated, the organization may face legal ramifications via lawsuits or governmental employment tribunals.

Network and Systems Security

Several policies provide standard guidelines for implementing network security principles within an organization and encompass areas such as the use of the Internet and internal network, data privacy, security incident response, human resources (HR) issues, and document security. These policies are often enforced by technical controls such as data loss prevention (DLP) tools that monitor for breaches of policy and issue a report when a breach occurs. Other tools may alert an administrator to machines joining the network that don't meet security requirements (having out-of-date antivirus signatures, for example) or report to an administrator when an unauthorized machine has been added to the network, or an inappropriate website has been visited.

There are a number of protection mechanisms on networks such as *firewalls* (which regulate network traffic), *intrusion prevention systems (IPSs)*, which protect traffic at a deeper level than firewalls, *security information and event management systems (SIEM)*, which aggregate log and event data from systems and alert personnel of threats), *threat intelligence platforms* (which assist a SIEM to recognize emerging threats), and so on. This is only a sampling of the protective technologies that are network-based. The deeper the understanding of the technologies, the more granular a risk assessor may be.

At the system or device level, many organizations publish a system and device *hardening* standard, a document that offers the principles of system hardening, as well as specific configurations required on all systems. The principles of system hardening include:

- Keep all systems and devices on the most current version of the operating system
- Install all required security and functionality patches

- Change all default login credentials
- Remove or disable all non-essential services and hardware
- Employ least privilege and need-to-know principles for access
- Forward all events to a centralized event log system, such as a SIEM
- Implement theft prevention measures, including whole-disk encryption

The philosophy one uses to set up the location of protective network devices can be referred to as the architecture. For example, for sensitive data, most compliance frameworks require that the data not be stored in the database on the same local network segment as a web server. This way, if the web server is compromised, attackers will not have immediate access to the data. If they reside on the same network, that may be considered a higher risk in a risk assessment—the recommendation being to move the database to another network behind the web server, with firewalls and intrusion prevention systems being the intermediary between the two local networks.

To monitor access and alterations to key configuration files and other files on a system, organizations employ *file integrity monitoring (FIM*—which detects changes to a file's content and access rights) and *file activity monitoring (FAM*—which detects access to sensitive files). When anomalies are detected, FIM and FAM tools send alerts to a SIEM or directly to system or security engineers.

Another factor to consider is *encryption* of data in motion. If a legacy system does not have encryption and it is sending data to another location on the Internet, there is a risk to anyone between the networks for *eavesdropping* or a *man-in-the-middle attack (MITM)*. A VPN tunnel can be set up between the two networks to ensure that at least the traffic is encrypted as it traverses the networks (or the Internet).

Vulnerability Management

Vulnerability management is a formal set of practices to detect exploitable vulnerabilities in networks, devices, systems, applications, facilities, and business processes. Vulnerability management depends on the accuracy and completeness of asset inventories, as well as change management, configuration management, and system and device hardening. Organizations lacking mature vulnerability management practices are far more likely to experience security incidents and breaches.

Vulnerability management is discussed fully in Chapter 2.

Security Incident Response

Security incident response is a formalized set of practices used to detect, analyze, contain, eradicate, and recover from security incidents while minimizing business impact and risk exposure. From a CRISC perspective, incident response is not merely a technical activity; it is a risk response capability that enables organizations to manage realized risk events, limit losses, and protect enterprise objectives.

A *security incident* is an adverse event that compromises, or has the potential to compromise, the confidentiality, integrity, or availability of information assets or systems. Examples include malware infections, unauthorized access, data exfiltration,

denial-of-service attacks, insider misuse, and third-party breaches. Not all security events rise to the level of an incident; incident response practices, therefore, rely on defined criteria and escalation thresholds to distinguish routine events from incidents requiring formal response.

An effective incident response practice is typically documented in an *incident response plan (IRP)*. The IRP defines roles, responsibilities, authorities, communication paths, decision-making processes, and required actions across the incident lifecycle. From a governance standpoint, the plan should clearly identify incident ownership, executive escalation points, legal and regulatory notification requirements, and integration with business continuity and disaster recovery processes.

Incident response activities are commonly organized into the following phases:

- **Preparation:** Preparation establishes the organizational capability to detect and respond to incidents effectively. This includes defining policies and procedures, assigning roles (such as incident commander and communications lead), training personnel, maintaining tools and logs, and conducting tabletop exercises. Preparation reduces risk by improving response speed, consistency, and decision quality.
- **Detection and analysis:** Detection involves identifying potential incidents through monitoring, alerts, reports, or third-party notifications. Analysis determines the nature, scope, severity, and potential business impact of the incident. Analysis may involve a *forensics* investigation, in which qualified forensic examiners look for clues to determine how the incident occurred. From a risk management perspective, this phase supports rapid assessment of likelihood and impact, enabling informed response decisions aligned with risk tolerance.
- **Containment:** Containment actions limit the spread or escalation of an incident. This may include isolating systems, disabling accounts, blocking network traffic, or suspending affected services. Containment strategies should balance technical risk reduction with business impact, recognizing that overly aggressive actions can introduce operational risk.
- **Eradication and recovery:** Eradication removes the root cause of the incident, such as malware, compromised accounts, or vulnerable configurations. Recovery restores systems and services to normal operation, often using backups or rebuilt environments. These activities directly support availability and resilience objectives and must be coordinated with business stakeholders.
- **Post-incident review:** After resolution, organizations should conduct a formal review to identify root causes, control gaps, and lessons learned. Findings should feed back into risk assessments, control improvements, training, and governance processes. This phase ensures that incidents contribute to continuous risk reduction rather than being treated as isolated events.

Security incident response practices must also address communication and reporting. Internal reporting ensures appropriate escalation to management and risk owners, while external reporting may be required to regulators, customers, partners, or law enforcement. Failure to manage communications effectively can significantly increase legal, regulatory, and reputational risk.



CRISC candidates need to understand that incident response is a corrective and adaptive risk response, which operates after a threat has materialized. IR plays a critical role in limiting impact, preserving evidence, supporting accountability, and strengthening the organization's overall risk posture.

Third-party Risk Management

Third-party risk management (TPRM) is a core function in information security programs in many organizations. Because so many organizations outsource large portions of their information processing, it's vital for organizations to fully understand and manage risks associated with each third party.

This topic is covered fully in Chapter 3.

Access Management

As an information security professional, you probably already know that a *security control* is a security measure or protection applied to data, systems, people, facilities, and other resources to guard them from adverse events. Security controls can be broken down and categorized in several ways. *Access controls* directly support the confidentiality and integrity goals of security, and indirectly support the goal of availability. Access control essentially means that we will proactively ensure that only authorized personnel can access data or the information systems that process that data. Access controls ensure that only authorized personnel can read, write to, modify, add to, or delete data. They also ensure that only the same authorized personnel can access the different information systems and equipment used to store, process, transmit, and receive sensitive data.

There are several different types of access controls, including identification, authentication, and authorization methods, encryption, object permissions, and more. Remember that access controls can be administrative, technical, or physical in nature. Administrative controls are implemented as policies, procedures, rules and regulations, and other types of directives or governance. For example, personnel policies are usually administrative access controls. Technical controls are those we most often associate with security professionals, such as firewalls, proxy servers, VPN concentrators, encryption techniques, file and folder permissions, and so on. Physical controls are those used to protect people, equipment, and facilities. Examples of physical controls include fences, closed-circuit television cameras, guards, locked doors, gates, and restricted areas.

In addition to classifying controls in terms of administrative, technical, and physical, we can also classify access controls in terms of their functions. These functions include preventative controls, detective controls, corrective or remedial controls, deterrent controls, and compensating controls. All the different controls can be classified as one or more of these different types of functions, depending on the context and the circumstances in which they are being used. Controls are described in more detail in Chapter 3.

The following access control principles and policies help provide a consistent organizational structure to prevent internal fraud and corruption in your organization:

- ***Least privilege*:** This principle grants users only the access rights they need to perform their job functions. This requires giving users the least amount of access possible to prevent them from abusing more powerful access rights.
- ***Separation of duties*:** This principle ensures that one single individual isn't tasked with high-security and high-risk responsibilities. Key critical responsibilities are separated between several users to prevent corruption.
- ***Job rotation*:** This provides improved security because no worker retains the same amount of access control for a position indefinitely. This prevents internal corruption by workers who might otherwise take advantage of their long-term position and security access.
- ***Mandatory vacation*:** A mandatory vacation policy requires employees to use their vacation days at specific times of the year or to use all their vacation days allotted for a single year. This policy helps detect security issues with employees, such as fraud or other internal hacking activities, because the anomalies might surface while the user is away. Increasingly, organizations are implementing a policy requiring mandatory administrative leave for situations in which an employee is under any sort of investigation, systems-related or otherwise.

Identification and Authentication

Identification and authentication are often misunderstood terms. They are related, to be sure, but they are not the same thing and really shouldn't be used interchangeably by a knowledgeable security professional. *Identification* refers to the act of an individual or entity presenting an identifier to a security system to assert that they are a specific entity. When you enter a username into a system, for example, or insert a debit card into an ATM, you are identifying yourself.

Authentication, on the other hand, is the second part of that process, where you present some form of proof of your identity, which is subsequently verified with a centralized database containing your authentication credentials. If the credentials you have presented match those in the authentication database, you are authenticated and allowed access to the network or resource. If they do not match, you are not authenticated and are denied access to the resource.

There are several different methods of identification and authentication, including single-factor (such as username and password, for example) and *multifactor authentication (MFA)*, which consists of two or more of the following: something you *know* (knowledge factor), something you *have* (possession factor), or something you *are* (biometric or inherence factor). Authentication also uses a wide variety of methods and technologies, such as Kerberos and 802.1X.

Authorization

Authentication to a resource doesn't automatically guarantee you have full, unrestricted access to a resource. Once you are authenticated, the system or resource defines what actions you are authorized to take on a resource as well as how you are allowed to interact with that resource. *Authorization* is what happens once you've successfully identified yourself

and have been authenticated to the network. Authorization dictates what you can and can't do on the network, in a system, or with a resource. This is usually where permissions, rights, and privileges come in. In keeping with the concept of *least privilege*, users should only be authorized to perform the minimum actions they need to fulfill their position's responsibilities.

Authorization has a few different components. First, there is *need-to-know*. This means that there must be a valid reason or need for an individual to access a resource, and to what degree. Second, an individual may have to be trusted, or *cleared*, to access a resource. This may be accomplished through a security clearance process or a non-disclosure agreement, for example.



Understand the differences between identification, authentication, and authorization. Remember that identification is simply presenting credentials, whereas authentication is verifying them. Authorization dictates what actions an individual can take on a system once authentication has occurred.

Accountability

Accountability means that a person is going to be held responsible for their actions on a system or regarding their interaction with data. Accountability is essentially the traceability of a particular action to a particular user.

Users must be held responsible for their actions, and there are different ways to do this, but it is usually ensured through *auditing*. First, there must be a unique identifier that is tied only to a particular user. This way, the identity of the user who performs an action or accesses a resource can be positively established. Second, auditing must be properly configured and implemented on the system or resource. What we are auditing is a user's actions on a system or interactions with a resource. For example, if a user named Sam deletes a file on a network share, we want to be able to positively identify which user performed that action, as well as the circumstances surrounding the action (such as the time, date, from which workstation, and so on). This can only be accomplished if we have auditing configured correctly and take the time to review the audit logs to establish accountability.



Although related, accountability is not the same thing as auditing. Accountability uses auditing as just one method to ensure that the actions of users can be traced to them and that they are held responsible for those actions. Other methods, such as non-repudiation, are used as well.

Non-repudiation

The concept of *non-repudiation* is closely related to accountability. Non-repudiation ensures that the user cannot later deny having performed an action, because the system is configured so that no one else could have performed it. The classic example of non-repudiation is given as the proper use of public key cryptography. If a user sends an email that is digitally signed

using their private key, they cannot later deny that they sent the email, since only they are supposed to have access to the private key. In this case, the user can be held accountable for sending the email, and non-repudiation is ensured.

Access Provisioning

Access provisioning is the act of configuring an access management system so that an employee, customer, or other subject is given access to a dataset, database, application, or system. Access provisioning is the final step in an access request process, a formal process in which access is requested, the request reviewed, and if approved, then a system configured to provide the requested access.

Because access management is a critical process and often functions as the only line of defense protecting a system from unauthorized access, organizations typically implement controls and business records associated with access management. If you examine any of the mainstream information security control frameworks (e.g., ISO/IEC 27001/27002, CIS Controls, PCI DSS, NIST SP 800-53), you'll see that each framework has numerous controls related to access management and access provisioning.

Many requests for assistance to an IT service desk are access-related. Thus, it's vital that organizations employ sound *identity proofing* procedures, which are steps to confirm the identity of the requestor. For critical systems such as financial accounting applications, IT departments may perform a segregation of duties (SOD) check, to see whether the access request, if approved, would create a segregation of duties problem. For instance, many accounting departments separate critical functions such as vendor creation, payment request, and payment approval to separate individuals as a measure of fraud prevention. If a user with payment request privileges is requesting payment approval privileges, the SOD check should flag this matter and recommend that the requested access not be granted.

Sometimes, an access request is of the form, "Please provide new employee Susan with all of the same access rights as existing employee Jennifer." If Susan and Jennifer have the same job title and responsibilities, this type of request, known as *access cloning*, may seem reasonable. The problem, however, is that Jennifer's access rights may exceed what is reasonable for a person in Jennifer's role; duplicating those rights for another user would increase risk all the more. For this reason, information security professionals discourage access cloning.

Access deprovisioning is the act of removing a user's or system's access to a dataset, database, application, or system. Access deprovisioning is performed when a user's role or responsibilities have changed or when a user's employment has been terminated, resulting in their no longer requiring access.

In termination matters, access deprovisioning becomes critically important, to prevent a terminated user from stealing or destroying information in an act of revenge. While information security standards often stipulate that a user's access should be removed within 24 (or 48) hours of termination, the best practice requires access to be removed just prior to the user being notified of the termination, thus preventing any emotionally fueled actions they would later regret.

Access Reviews

An *access review* is an activity where an approver reviews individual users' access rights to a system. The purpose of an access review is to confirm two things:

- That each subject continues to need access to specific roles, data, or systems
- That access provisioning and deprovisioning processes are working correctly

During an access review, an approver will be reviewing a manually created list, or interacting with an access management system. Either way, the reviewer considers each user to determine whether they still need access.

Access reviews are typically performed quarterly for critical systems, and somewhat less frequently for other systems, but generally they are performed at least annually. Access reviews for privileged users should be performed at least once per quarter, but monthly is preferred.

Access reviews are considered a corrective control to curb *privilege creep*, in which users gradually accumulate access rights over many years. As users transition from role to role within an organization, they are likely to acquire new access rights, but the old access rights they no longer need are often retained.

Privileged Access Management

Privileged access management (PAM) refers to the same access provisioning and access review processes for privileged users such as system administrators, domain administrators, network administrators, and other roles where users have administrative responsibilities for one or more devices, systems, or applications. Administrative responsibilities include the ability to change a system's or device's operating parameters, start or stop the device or system, and create new users for the device or system.

Larger organizations often employ PAM tools that manage privileged access, including storing privileged login credentials, brokering administrative access to systems, and even recording the actions that administrators take when logged in to systems with privileged credentials.

Many organizations practice *administrative account separation*, a form of least privilege access, whereby those personnel with privileged access to a system will also have non-privileged access to the system. For example, Chris, a domain administrator, has domain administrator login credentials to be used when Chris is performing administrative duties. However, when Chris is performing ordinary non-privileged duties, Chris logs in with separate, non-privileged credentials. One example of a good reason for administrative account separation is to reduce the impact of a malware attack. If Chris executes malware as a domain administrator, the impact of the malware may include many or even all systems in the organization. Other good reasons include preventing the possibility of inadvertently performing actions that may damage the infrastructure while logged in using a privileged account, as well as ensuring accountability for privileged users.

Physical Access Security

As part of its overall access control, your organization must have strong physical access controls and an associated policy, and it must ensure that all workers are educated on the policy's use. Depending on the security level of the organization, physical security may

include guarded or unguarded entrances. Even on guarded premises, the use of security access cards ensures that only identified, authenticated, and authorized persons can enter a facility. Each security access card is coded with the authorization level of the user, who will be able to access only areas of the facility required by their job function. For example, only network and systems administrators would be able to access a server and network communications room with their access card.

Workers must be trained to always close automatically locking doors behind them and not allow unidentified people to follow them through (known as *tailgating*). Most security access cards include a photograph of the cardholder to further identify users in the event they are challenged for their identity. Workers must be encouraged to report suspicious individuals within the premises who are unfamiliar and do not have proper identification.

A published organizational security policy that addresses physical access allows your workers to have proper knowledge of security procedures and to be equally active in the responsibility for physical security.

The Role of HR

An organization's human resources department plays a vital role in identity and access management. HR is usually the owner of record of the list of employees in an organization, together with their job title and placement in the organization chart. Identity and access management processes rely on the accuracy and integrity of the organization's list of employees: when reviewing an access request, the subject's employment, job title, and org chart placement should be reviewed, to ensure that the subject still works in the organization and that their access request aligns with their title and role. Any errors or latency in HR's work of keeping the list of employees up to date will increase access-related risk.

HR departments also typically manage:

- **The recruiting and hiring** process, including conducting *background checks* and ensuring a candidate has the legal right to work in the organization.
- **Performance management,** including performance reviews and disciplinary action.
- **Training and training records,** and often the management of a learning management system (LMS) to electronically deliver training content and maintain records of attendance and completion.
- **Collecting signatures,** including documents required by local governments, acknowledgments of corporate policies, non-compete agreements, and *non-disclosure agreements (NDAs)*.



What About Contractors, Consultants, and Vendors?

In many organizations, HR plays no role in the management of contractor, consultant, and vendor personnel. Hence, access requests, reviews, and terminations related to contractors, consultants, and vendor personnel may not be able to rely on HR's accurate and up-to-date list of current personnel. Organizations with numerous contractors and consultants will need to find another way of tracking those personnel, to avoid undue access-related risks.

Security Architecture

Organizations use different means for implementing a *security architecture*. Because security architecture is both a big-picture pursuit and a detailed plan, different decision-makers may be involved, and different business processes are used.

At the enterprise level, the organization will use its policy and governance functions to make decisions about significant aspects of security architecture, such as decisions about the brands of servers, workstations, and network devices to be used.

At the detail level, processes such as configuration management and change management are used to implement device characteristics on individual devices or groups of devices. For instance, an upgrade to an organization's Domain Name System (DNS) infrastructure might result in an increase in the number of name servers in use, necessitating that most or all devices in the organization be updated accordingly.

Over time, organizations will make changes to their architecture models (both at the macro level and micro level). For instance, changes and advances in technologies such as software-defined networking (SDN), virtualization, and microservices architectures will compel organizations to revise their enterprise security architecture. However, updating architecture documents is only the beginning. Organizations will also initiate tasks or projects to implement these improvements in their environment. Some of these changes, such as updates to industry device-hardening standards, may result in little more than change management to update a few configuration settings, but other changes may involve multiyear projects to modernize some aspect of an organization's environment.

The *Sherwood Applied Business Security Architecture (SABSA)* model is used in the development of an information security architecture. SABSA is similar in its structure to the Zachman Enterprise Architecture Framework.

Information Security Awareness Training

Both security awareness and risk awareness training programs are critical controls to reduce risk within an organization. Training can cover a whole range of topics, including phishing, data handling, data sensitivity, compliance requirements, wireless challenges, how data is lost, how to submit incidents, and so on. Both help improve the risk posture of organizations at an individual level by increasing general employee awareness of threats and required practices, as well as leader awareness of organizational risks. Not having any training is an indication, from a risk perspective, that people may not know how to handle different types of risks or even how to recognize risks.

Awareness programs are a very necessary part of risk management. They can't be viewed as simply another two-hour training session to check a box for management or compliance. Awareness is essential because it helps form and maintain the organization's risk culture. It also educates personnel at all levels of the organization, including employees and other workers, managers, and senior leadership, on the organization's risk strategy, its appetite

and tolerance levels for risk, its risk management plan, and other relevant topics necessary to manage risk in the organization. Beyond the education on organizational governance and risk management processes, awareness training can give all members of the organization the knowledge they need to better identify, assess and evaluate, and respond to risk. Awareness training may be required for compliance with governance in some cases, but even if it's not, it should be considered critical to the overall risk management strategy in the organization and given due consideration in the organizational priority list.

Establishing the security and risk awareness programs in an organization can be a challenge. One way organizations fail is to simply direct someone to develop a training program, when the organization has not even established its risk management strategy or plans. Developing the organization's approach to risk is an essential first step before implementing awareness training. The organization must develop, formally if possible, its stance on risk appetite and tolerance as well as its risk management strategy. It should decide what risk management methodologies it will use, as well as what standards and frameworks. Only then can a training program be developed based on a good, solid risk management framework within the organization.

Additionally, the training program should be periodically evaluated for effectiveness, as well as updated with current risks, governance, tools, and techniques.

The following two sections discuss the different tools and techniques used in awareness training and how to develop an awareness program within an organization.

Designing Training Programs

As with any program, it's crucial that an organization identify the core objectives of a security awareness program. Organizations often implement a security awareness program to comply with a specific regulation or framework, but sometimes a security awareness program is implemented to reduce risk. The latter often requires more rigor and depth in security awareness training, as "checkbox compliance" is characterized by performing the least amount of effort required to comply with a directive.

Audience-specific Content

Information security and risk affect virtually all workers in an organization. Because personnel across an organization have a wide breadth of skills and knowledge, security awareness training often includes two components:

- **General training for all workers:** This training generally includes information about the organization's acceptable use policy and other policies, corporate use of information systems, the use of personally owned devices to conduct company business, email and Internet hygiene, and more.
- **Specific training for technical workers:** This training targets a variety of technical users, including developers, engineers, and specialists in areas such as privacy and AI. Training content will include any technology-specific policies and standards, plus in-depth content on role-specifics such as safe software coding, system hardening standards and procedures, or privacy policies and practices.

Behavioral Reinforcement

Better security awareness training programs provide opportunities for employees to practice their skills. Workers who participate in awareness training should be able to easily put into practice the concepts, tools, and techniques they learn to better protect critical information and systems. Some examples include:

- **Password strength:** Recognizing and creating strong passwords.
- **Phishing:** Recognizing phishing messages and acting appropriately.
- **Social engineering:** Recognizing common social engineering situations, such as attempting to get assistance from the IT service desk and tailgating.
- **Rewards and recognition:** Employees who report actual attack scenarios may be rewarded with prizes, encouraging others to be more vigilant.

Leadership-level Risk Training

Management should be committed to awareness training as part of its overall risk management strategy. The training program should articulate not only the risk culture of the organization but also the different risk management needs the organization has. These might include specific risk factors, threats, vulnerabilities, and so on, considered when the training is developed.

Beyond initial or recurring awareness training, ongoing communication within the organization is a must for effectively managing both security and risk awareness. Workers in general, but also specifically those with key risk management duties, should be given information on an ongoing basis regarding organizational risks and how to manage and deal with them. Obviously, some information would be restricted from the general organizational population, but specific instances of threats and vulnerabilities, risk factors, and so on should be provided to risk managers in key areas so they can keep updated on the most current risk posture for the organization.



While training programs are sometimes the first things that are cut from the budget or the last things to be developed in a program, don't underestimate the importance of awareness training within the overall risk management strategy. Training not only can make the risk management process within the organization more effective, but it can also help reduce or mitigate risk because it educates people about the subject, and this alone may minimize risk.

Training Delivery and Measurement

Like most training, awareness training should meet several criteria. First, it should be geared toward specific groups of audiences. Second, training shouldn't be a one-time event. Typical training ran from 30 to 120 minutes and longer. While this approach may be considered efficient, it is no longer considered effective. A preferred approach is to

instead offer training far more often, perhaps monthly, with shorter training sessions. Organizations can even employ microtraining sessions of just a few minutes each, offered weekly. Periodic, recurring training is a good idea simply because it can be used to reinforce and refresh stale knowledge and bring workers up to date on the latest tools, techniques, and risk considerations. Finally, awareness training should be well organized and conducted by knowledgeable instructors, both from inside and outside the organization. Internal trainers can give the benefit of the organization's specific views on risk culture, appetite, and tolerance, while external trainers bring the benefit of objective knowledge and risk management methods from the industry.

An organization can deliver awareness training in several different ways, and using various combinations of training types can help deliver an effective training program. Classroom training is, of course, one standard method. Other methods might include individual-based training that comes from reading, computer-based training, and so on. Workers might also be required to read a risk management handbook that defines the different rules and regulations covering risk within the organization. Specialized training on risk management might have to be provided by an external training provider for those individuals with defined risk management responsibilities.

Because security awareness training is often required by specific regulations or standards, organizations are implicitly obligated to maintain accurate training records, permitting the organization to objectively demonstrate compliance. Beyond compliance, many organizations include competency testing as a part of security training, and record test scores. These records support non-repudiation, meaning that an employee who knowingly violates policy cannot claim they weren't aware, since their test scores show they were. Beyond compliance, metrics and KPIs for security awareness training can demonstrate risk levels and even serve as a leading indicator for future events.

Business Continuity Management

Business continuity management is the high-level, overall effort to ensure that an organization's critical processes and functions will be able to continue operations despite disruptive incidents, such as natural disasters like hurricanes and earthquakes, as well as human-caused disasters like utility interruptions and social unrest.

BCM Foundations

Business continuity management practices follow a top-down, industry-standard methodology that consists of these steps:

- **Step 1: Define scope.** Executive management defines the outer boundaries for BCM, usually aligned to business units or geography.
- **Step 2: Conduct a business impact analysis.** This is a review of business processes to identify the most critical processes in the organization, together with their dependencies.

- **Step 3: Establish key recovery objectives.** These parameters drive the speed and completeness of recovery after a disaster.
- **Step 4: Develop business continuity plans.** These are procedures to be followed in the event of a disruption.
- **Step 5: Test business continuity plans.** These tests ensure personnel are familiar with procedures, and that those procedures will actually work in a real emergency.
- **Step 6: Develop disaster recovery plans.** These are procedures to be followed by an IT department to recover primary information processing systems or move information processing to alternate systems. This is covered in Chapter 4.
- **Step 7: Test disaster recovery plans.** These tests ensure personnel are familiar with procedures, and that those procedures will actually work in a real emergency. See Chapter 4.
- **Step 8: Execute business continuity and/or disaster recovery plans when a disruption occurs.**
- **Step 9: Perform an after-action review.** This is done to identify successes and improvement opportunities after a disruption event has ended.

Business continuity planning involves several steps. Since the focus of this domain isn't solely on business continuity and disaster recovery, I'm not going to go into an in-depth treatment of these processes. However, I do discuss some of the basic concepts of business continuity and disaster recovery as they relate to risk. Two of the key processes involved in business continuity planning are conducting a *business impact analysis (BIA)* and determining a recovery strategy.

Business Impact Analysis

A business impact analysis (BIA) is used to identify the most critical business processes in an organization, together with the dependencies for each process. For each business process, the BIA captures the impact on the organization should the process be impaired by a disaster. A BIA closely resembles some of the steps you would take during a risk assessment, simply because you are identifying assets (which include, of course, people, equipment, facilities, and processes) and prioritizing those assets. The reason they are prioritized is that in the event a disaster were to occur, the business would have to rely on specific assets or processes more so than others to resume critical processes. These critical assets or processes should be brought back online and recovered first, but to accomplish this, they must be identified by the organization, and plans must be made for their recovery. That's where the business impact assessment comes in. Business impact analysis is described in more detail in Chapter 2.

Recovery Objectives

Two key areas of concern in business continuity are the *recovery point objective (RPO)* and the *recovery time objective (RTO)*. But before RPO and RTO are established, management must first determine *maximum tolerable downtime (MTD)*, which is the length of time that an outage would jeopardize the ongoing viability of the organization.

Once MTD is set and approved by management, it is possible to go back and establish RPO and RTO for each business process. Both areas can be measured in the context of time. The recovery point objective is the amount of data that can be lost based on a time measurement. In other words, if the recovery point objective is two days, the maximum amount of data that can be lost without a profound effect on the organization's ability to function would be two days' worth of data. Note that this measurement isn't concerned with the actual amount of data in terms of quantity, such as gigabytes, for example; instead, it's concerned with measuring how many days or hours' worth of data can be lost, at most, without impeding the function of the business. The recovery time objective, also measured in terms of time, is the maximum amount of time the organization can afford to lose before it starts recovery. This measurement of time could be days, hours, minutes, or even seconds.



Recovery objectives, including RPO and RTO, influence the development of business continuity plans as well as disaster recovery plans.

Resilience Engineering

Resilience is a term that refers to the ability of an organization to survive disruptive events and continue with its mission and function. In other words, it could be said that resilience is the ability of an organization to resist the effects of adverse events and its ability to bounce back after one of these events. Resilience can affect the degree to which different risk factors, as well as threats and vulnerabilities, have on impacting the normal business operations of the organization. A more resilient organization can better tolerate risk factors and the effects they have on the organization's ability to prepare for and respond to adverse events.

Various risk factors can affect the organization's ability to respond to disasters as well as to continue the business and mission functions after it has recovered from a disaster. Some of these risk factors are inherent to the organization itself; these include the size of the organization, its facilities, its people, and how well it has performed business continuity and disaster recovery planning. A higher level of management commitment to business continuity planning reduces the effect of risk factors on the business; less commitment may render the organization more susceptible to these risk factors. A larger organization may have more complexity in terms of the number of people to take care of during a disaster, as well as the number of critical assets or processes that must be planned for. Small organizations don't necessarily have an advantage; however, a small organization may have fewer resources to adequately plan for and address business continuity and disaster recovery.

Another risk factor that an organization could face is physical location and geography. For example, a business situated on the U.S. Gulf Coast is probably more susceptible to adverse weather, such as hurricanes and tornadoes, than a business located inland. Businesses that are in areas prone to heavy flooding, or have forest fires or extreme snowstorms at various times of the year, should plan for these weather conditions appropriately. Distance could also be a risk factor; if the organization is located farther away from population centers or infrastructure services (such as fire departments, police, hospitals,

telecommunications providers, and so on), there is a greater risk that the organization could not adequately respond to a disaster.

There are also risks associated with third-party providers, and these providers can affect business continuity and disaster recovery as well. Yet another risk factor could be the degree of outsourcing the organization is involved in. A business that outsources some of its processes or key functions may have less control over those functions or processes during an incident or disaster. The business may be at the mercy of the third party with whom they are contracted and subject to the third party's disaster recovery strategies (or lack thereof). Business continuity and disaster recovery are two areas that should be carefully considered when engaging in business with any external provider, usually through a service-level agreement or another method.

In addition to the factors I have already mentioned, the very nature of the organization's mission or business could be a risk factor. Organizations that have critical processes or produce critical goods and services that are time-sensitive, for example, may incur more risk than an organization that can afford to be down or out of service for a few days. Time sensitivity can increase the impact of a disaster on an organization, so it should be considered when developing continuity and recovery strategies. Market factors, such as the perishability of goods, for example, can become risk factors that an organization must deal with as well.

Most people associate the threats involved with business continuity and disaster recovery with natural disasters. Obviously, natural threats exist: severe weather, earthquakes, fires, and so on. These are threats that can't be typically prevented or neutralized, but they can be planned for. Hardened facilities, adequate ventilation, fire suppression systems, and other physical or environmental controls can mitigate the damage from these threats. On the other hand, human-made threats can also affect the ability of an organization to recover from an incident or continue to function. These can be broken down into unintentional and intentional threats. Intentional threats include those associated with hackers, terrorists, theft, and other malicious actors and activities. Unintentional threats may be attributed to accidents that are caused by carelessness, lack of training, and so on. All of these threats, both natural and human-made, are usually the types that cause incidents or disasters. However, there are also threats to the continuity management process itself.

Threats to the business's ability to manage its continuity and disaster recovery include incomplete or inadequate planning as well as limited resources to carry out any BCP/DRP activities. Because of regulatory considerations in many cases, a business usually must have some degree of business continuity planning accomplished; the real threat centers on whether it is enough and adequately covers the different facets of recovery and continuity that the business must consider. Usually, inadequate planning and insufficient resources committed to this area are a result of shortsightedness on the part of the organization's leadership and management.

A wide variety of vulnerabilities can affect the ability of an organization to respond to and recover from an incident or disaster. The same vulnerabilities can also, in turn, affect the organization's ability to continue to function at some level or degree after the disaster. Some of these vulnerabilities involve the lack of resources, such as a lack of personnel, equipment, facilities, or dedicated recovery resources. For instance, if the organization has not acquired an alternate processing site, there may be no way it can quickly recover its operations to a functional state. Other vulnerabilities include the lack of training for personnel tasked with disaster recovery or business continuity responsibilities.

Additionally, there are also process vulnerabilities. Process vulnerabilities are those that result from a failure of the organization's management to adequately plan for, staff, and equip continuity management efforts. For example, a lack of management commitment is a vulnerability that will undoubtedly lead to a lack of training, resources, and personnel dedicated to continuity management.

An inadequate or nonexistent business impact assessment is a vulnerability because an organization won't be able to adequately identify its critical processes or assets and determine how the loss of those critical items impacts the organization. Another critical vulnerability could be faced if the organization has failed to adequately test the response plan.



Be familiar with the different business process areas and their respective risk factors, threats, and vulnerabilities. Remember that risk factors, while not a threat or a vulnerability, contribute to the overall risk of the organization by increasing or decreasing the likelihood or impact of threats exploiting vulnerabilities or the organization's resilience to disruptive events.

Data Privacy and Data Protection Principles

Throughout this chapter, I have discussed the core concepts of information security that any program should be striving for. This section outlines key data privacy and data protection principles and associated policies and procedures that organizations should consider when aiming to reduce risk.

Privacy policies are agreements that protect individually identifiable information. An organization engaged in online activities or e-commerce has a responsibility to adopt and implement a policy to protect the *privacy* of personally identifiable information (PII). Increasingly, regulations such as the European Union's General Data Protection Regulation (GDPR), the California Consumer Privacy Act (CCPA), the California Privacy Rights Act (CPRA), and the U.S. Health Insurance Portability and Accountability Act (HIPAA) require a privacy policy that is acknowledged before use. Organizations should also take steps to ensure online privacy when interacting with other organizations, such as business partners. Privacy obligations also extend to an organization's use of the PII of its workers.

The following recommendations pertain to implementing privacy policies:

- An organization's privacy policy must be easy to find, read, and understand, and it must be available prior to or at the time the PII is collected or requested.
- The policy needs to state clearly what information is being collected; the purpose for which that information is being collected; possible third-party distribution of that information; the choices available to an individual regarding the collection, use, and distribution of the collected information; a statement of the organization's commitment to data security; and what steps the organization takes to ensure data quality and access.

- The policy should disclose the consequences, if any, of a person's refusal to provide information or the refusal to permit its processing.
- The policy should include a clear statement of what accountability mechanism the organization uses, such as procedures for dealing with privacy breaches, including how to contact the organization and register complaints.
- Organizations must provide an opportunity for data subjects to have an opportunity to exercise choice regarding how PII collected from them online can be used when such use is unrelated to the purpose for which the information was collected. At a minimum, organizations should provide data subjects the opportunity to opt out of such use.
- When an individual's information collected online is to be shared with a third party, especially when such distribution is unrelated to the purpose for which the information was initially collected, the individual should be given the opportunity to opt out.
- Organizations creating, maintaining, using, or disseminating PII should take appropriate measures to ensure its reliability and should take reasonable precautions to protect the information from loss, misuse, or alteration.

Each organization must evaluate its use of the Internet and its internal systems to determine the type of privacy policy it needs to protect all involved parties. The privacy policy helps protect the organization from legal issues, raising employees, customers, and constituents' comfort levels regarding the protection of their information. A privacy policy should include the following:

- **Information collection:** The organization will collect, use, and exchange only data pertinent to the exact purpose, in an open and ethical manner. The information collected for one purpose shouldn't be used for another.
- **Direct marketing:** The organization can use only non-PII for marketing purposes and must certify that the persons' personal information will not be resold to third-party organizations.
- **Information accuracy:** Ensure the data is accurate, timely, complete, and has been collected in a legal and fair manner. Allow people the right to access, verify, and change their information in a timely, straightforward manner. Inform people of the data sources and allow them the option of removing their names from the marketing lists.
- **Information security:** Apply security measures to safeguard the data on databases. Establish worker training programs and policies on the proper handling of PII. Limit access to a need-to-know basis on personal information and divide the information so that no one worker or unit has the whole picture. Follow all government regulations concerning data handling and privacy.



Privacy policies must be easy to find, and they must provide information on how to opt out of any use of personally identifiable information.

Data Sensitivity and Classification

An *asset* is a general, all-encompassing term that could include anything of value to an organization. The term *asset* can be applied to data, systems, capabilities, people, equipment, facilities, processes, proprietary methods, and so on; it is anything the organization values and desires to protect. Organizations typically determine how important their assets are to them and how much protection should be afforded to those assets. For example, intellectual property is an extremely valuable asset to the organization and is usually well-protected. This is really the fundamental concept of risk management: how much security or protection a particular system or piece of data requires, based on how likely it is that something bad will happen to it, balanced with what the organization can really afford to spend on the protection for that asset. To make reasonable decisions on how much security an asset needs, the organization must decide how much the asset is worth. I discuss worth in terms of dollars a bit later in the chapter, but for now, let's look at it from the perspective of asset sensitivity. In terms of sensitivity, you'll usually see the term *data sensitivity* in particular, but we could also broadly apply sensitivity to any asset in an organization.

Data (or other asset) sensitivity refers to how much protection the organization feels a particular system or piece of data requires, based on its value to the organization and the impact if it were lost, stolen, or destroyed. For example, information published on the organization's public website or in the company newsletter is public knowledge and is usually easily retrievable if, for some reason, the hard disk containing that data fails or is erased. Since the data is public, we may not consider that data to be very sensitive in nature and therefore require less protection for it than other types of data. On the other hand, customer order data is critical to the organization, simply because its business operations depend on that data to function and turn a profit, so it makes reasonable sense that the organization would spend a little bit more time, money, and effort in protecting that particular data. Therefore, its sensitivity, or *classification level*, would be considered somewhat higher than public data. Generally, the higher the sensitivity of the data, the more protection it is given.

In basic security classes, you typically learn about the different classifications of data found in both commercial organizations and government agencies. In commercial organizations, typical data sensitivity labels include Private, Company Sensitive, Proprietary, and so on. In the U.S. government, data sensitivity levels include Confidential, Secret, and Top Secret, and are classified based on the level of damage to the security of the United States that could be incurred if data at these various classification levels were disclosed or lost. Remember that data sensitivity is driven by the value of the data to the organization, the impact if it is lost, stolen, or destroyed, and balanced by the commitment of resources the organization is willing to provide to protect that data. Data sensitivity and classification policies specify the different formal levels of sensitivity in the organization and what those levels require in terms of protection.

Data Protection

Information security is ultimately about protecting sensitive information. Data protection begins with a data classification policy (discussed above) with accompanying data handling guidelines that stipulate how data at various classification levels and use cases are to be protected.

Because data exists in many forms and is used in many ways, there are many methods used to protect data. These methods include:

- **Access rights:** Discussed in the "Access Management" section earlier in this section, access to data in its various forms (flat files, databases, binary large objects) is granted only to those personnel with a verified need.
- **Backup/replication:** Discussed in the "Disaster Recovery Management" section in Chapter 4, organizations employ backup, replication, and other means to ensure that critical data can be recovered as needed.
- **End-user behavior:** Detective tools, including *end-user behavioral analytics (EUBA)*, learn end-user (and system) normal behavior, and generate alerts when anomalous activities are seen.
- **Data loss prevention:** *Data loss prevention (DLP)* tools reside in systems and on networks and are used to detect and (optionally) block the storage and movement of data that violates data classification policy. DLP systems are used to detect instances where users and attackers are moving data out of the organization (known as *exfiltration*) in violation of data classification policy and guidelines.
- **Social media monitoring:** While only authorized personnel are permitted to represent an organization on social media, organizations employ monitoring tools to look for unhappy customers, incidents, and employees posting information about the organization they ought not do.
- **Data governance:** Data governance exists as an overall management activity to monitor data management and operations, informing management and executives of the data that exists in the organization, where it resides, how it is used, and how long it should be retained.

Privacy as a Risk Scenario

From a risk management perspective, privacy should be understood not merely as a compliance obligation, but as a source of distinct and material risk scenarios that can affect an organization's strategic objectives, operations, and reputation. Privacy risk arises when personal data is collected, processed, stored, shared, or retained in ways that create the potential for regulatory non-compliance, unauthorized disclosure, misuse, or loss of trust.

A privacy risk scenario typically involves a combination of a threat event, a vulnerability in data handling practices, and an adverse impact or harm. Common threat events include cyberattacks targeting personal data, insider misuse, third-party mishandling of data, system misconfigurations, and failures in consent or access controls. Vulnerabilities may stem from excessive data collection, unclear data ownership, weak data classification, inadequate retention controls, insufficient vendor oversight, or ineffective incident response procedures. When these threats exploit such vulnerabilities, the resulting impact can be significant. Just read today's news, which is likely to include a story on a new privacy breach, and legal/investigative updates on privacy breaches already known.

The impact of privacy risk events on organizations extends beyond immediate technical consequences. Regulatory penalties and enforcement actions may result from violations of applicable privacy laws or contractual obligations. Many regulations impose mandatory

breach notification requirements, which can increase reputational damage and customer attrition. Legal exposure, including civil litigation and class-action lawsuits, is also a common outcome. In addition, privacy incidents often disrupt operations, consume management attention, and erode trust among customers, employees, partners, and regulators.

Privacy risks are further complicated by third-party vendor relationships. Organizations often rely on vendors, cloud providers, processors, and service partners to store and process personal data. If third parties fail to implement adequate safeguards or exceed their authorized use of data, the originating organization may still bear regulatory and reputational responsibility. As a result, privacy risk scenarios must explicitly account for data sharing, outsourcing, and cross-border processing arrangements.

Effective management of privacy risk requires integrating privacy considerations into the broader risk management lifecycle. During risk identification, organizations should inventory personal data, understand processing purposes, and identify applicable regulatory obligations. Risk analysis should evaluate both the likelihood and impact of privacy-related events, considering factors such as data sensitivity, volume, and exposure. Risk response strategies may include implementing technical and organizational controls, reducing data collection, strengthening vendor oversight, revising retention schedules, or formally accepting residual risk within approved tolerance levels.

It is essential to recognize that privacy risk is ongoing and dynamic. Changes in business models, technology platforms, regulatory expectations, or data usage practices can introduce new privacy risks or alter existing ones. Treating privacy as a set of realistic risk scenarios—rather than as static policy compliance—supports more effective governance, informed decision-making, and sustained risk reduction across the organization.

Privacy Regulatory Foundations

Privacy regulations concerning the electronic collection and use of sensitive information have been in existence since the 1990s. As information processing practices have evolved, privacy-related laws have likewise been updated to counter emerging threats and concerns. Regulations you should be aware of include the GDPR, EU AI Act, CCPA/CPRA, and other sectoral laws.

GDPR

The *General Data Protection Regulation (GDPR)* is a comprehensive data protection regulation enacted by the European Union and has been in effect since 2018. Its purpose is to protect the personal data and privacy of individuals located in the EU and to harmonize data protection requirements across member states. GDPR applies not only to organizations established in the EU, but also to organizations outside the EU that process personal data of individuals in the EU, making its scope extraterritorial.

GDPR is principles-based and centers on lawful, fair, and transparent processing of personal data. It establishes six core data processing principles:

- Lawfulness, fairness, and transparency
- Purpose limitation

- Data minimization
- Accuracy
- Storage limitation
- Integrity and confidentiality

Organizations must be able to demonstrate compliance with these principles.

GDPR grants enforceable rights to *data subjects*, including the right to access, rectify, erase (the well-known "right to be forgotten"), restrict processing, object to processing, and receive their data in a portable format. Organizations must implement processes to respond to these requests within defined timeframes, typically one month.

From a risk management perspective, GDPR introduces explicit obligations that increase regulatory, legal, and reputational risk. Organizations are required to implement appropriate technical and organizational measures to protect personal data, conduct data protection impact assessments (DPIAs) for high-risk processing activities, and notify supervisory authorities of certain personal data breaches within 72 hours. In some cases, affected individuals must also be directly notified.

GDPR enforcement includes significant administrative fines, which may reach the greater of €20 million or 4% of global annual turnover. For CRISC candidates, GDPR should be understood as a source of mandatory risk treatment requirements that influence governance, control design, incident response, third-party risk management, and ongoing compliance monitoring within an organization's broader risk management program.

EU AI Act

The *EU AI Act* is a comprehensive regulatory framework governing the development, deployment, and use of artificial intelligence systems within the European Union. While not a privacy law per se, the Act has significant privacy and data protection implications. The Act adopts a risk-based approach, classifying AI systems into prohibited, high-risk, limited-risk, and minimal-risk categories. High-risk AI systems are subject to strict requirements related to data governance, including controls over training, validation, and testing data to reduce bias, ensure data quality, and protect fundamental rights.

From a risk management perspective, the EU AI Act introduces new compliance, governance, and accountability risks, particularly where AI systems process personal data or make automated decisions affecting individuals, including employees, customers, and constituents. Organizations must consider the interaction between the EU AI Act and existing privacy laws, such as GDPR, when assessing AI-related privacy risk, third-party dependencies, transparency obligations, and incident response requirements.

CCPA/CPRA

The *California Consumer Privacy Act (CCPA)*, effective January 2020, and its amendment, the *California Privacy Rights Act (CPRA)*, effective January 2023, establish comprehensive data privacy requirements for residents of California. Together, they provide data subjects with rights over the use of their personal information and impose obligations on organizations that collect, use, sell, share, or retain that information. These laws apply to

certain for-profit entities that do business in California and meet defined thresholds related to revenue, data volume, or data monetization.

CCPA and CPRA grant data subjects specific rights, including the right to know what personal information is collected and how it is used, the right to access and delete personal information, and the right to opt out of the sale or sharing of personal information. CPRA expands these rights by introducing the right to correct inaccurate personal information and by creating enhanced protections for sensitive personal information, such as precise geolocation, government identifiers, and financial data.

From a governance and risk perspective, CPRA significantly strengthens enforcement and accountability. It establishes the California Privacy Protection Agency (CPPA) as an independent regulator and introduces explicit requirements for data minimization, purpose limitation, and retention limits. Organizations must implement reasonable security procedures and document privacy practices through notices, contracts with service providers, and internal controls.

CCPA/CPRA also increases compliance risk by expanding breach-related liability and enforcement authority. Statutory penalties may be imposed for non-compliance, and individuals may bring private actions in the event of certain data breaches. For CRISC candidates, CCPA and CPRA should be understood as regulatory drivers of privacy risk that influence data governance, third-party risk management, incident response, and ongoing compliance monitoring within an enterprise risk management framework.

Sectoral Laws

In addition to comprehensive privacy regulations, many jurisdictions impose sector-specific laws that include privacy and data protection requirements. These laws apply to particular industries and create mandatory obligations that must be considered in risk assessments, control design, and compliance monitoring. These regulations include:

- **Health Insurance Portability and Accountability Act (HIPAA):** HIPAA governs the use and disclosure of protected health information (PHI) by covered entities and their business associates in the United States. It establishes privacy, security, and breach notification requirements, mandating administrative, technical, and physical safeguards to protect PHI. From a risk perspective, HIPAA introduces regulatory, legal, and reputational risk tied to healthcare data handling and third-party relationships.
- **Gramm-Leach-Bliley Act (GLBA):** GLBA applies to financial institutions and requires the protection of consumers' nonpublic personal information (NPI). The Safeguards Rule mandates the implementation of a written information security program, while the Privacy Rule governs disclosure of customer information. GLBA creates ongoing compliance and oversight obligations that directly influence information security governance and vendor risk management.
- **Payment Card Industry Data Security Standard (PCI DSS):** While not a law, PCI DSS is an industry-mandated standard with strong privacy and data protection components for payment card data. Non-compliance can result in fines, loss of card-processing privileges, and increased fraud risk, making it relevant from a risk management standpoint. More details about PCI DSS are found in Chapter 3.

- ***Family Educational Rights and Privacy Act (FERPA)*:** FERPA protects the privacy of student education records in U.S. educational institutions. It limits disclosure of personally identifiable information and grants rights to students and parents, introducing compliance and reputational risk for covered organizations.
- **U.S. State Breach Notification Laws:** Many U.S. states have laws requiring notification to individuals and regulators following certain personal data breaches. These laws increase incident response complexity and elevate legal and reputational risk.

CRISC candidates should understand sectoral laws as mandatory risk drivers that impose specific control, reporting, and governance requirements based on industry context.



Risk managers must understand the highly dynamic nature of information security, privacy, and AI-related regulations, and employ methods to stay abreast of emerging laws and other legal developments.

Summary

The three goals of security, known as the CIA triad, are confidentiality, integrity, and availability. Supporting these three goals are other security elements, such as access control, data sensitivity and classification, identification, authentication, authorization, accountability, and, finally, non-repudiation.

Standards, frameworks, and practices relevant to risk management include NIST RMF, COBIT 2019, and the various ISO standards. Industry frameworks, standards, and practices can be adopted when building an information security program, rather than building these elements from scratch.

The NIST Risk Management Framework (RMF) is a seven-step methodology that provides for risk management all the way through the information systems lifecycle. The NIST Cybersecurity Framework (CSF) is a framework for managing and reducing cybersecurity risk. Various ISO/IEC standards establish requirements and frameworks for information security governance and risk management, including ISO/IEC 27001, ISO/IEC 27002, ISO/IEC 27701, and ISO 31000. COBIT is a management framework developed by ISACA that facilitates the governance of enterprise information and technology. COBIT consists of two layers, governance and management, which are further broken down into a total of 40 separate objectives. The Risk IT Framework draws on traditional risk management principles from various enterprise risk management (ERM) standards and describes activities and processes considered best practices.

Security policies, standards, guidelines, and procedures are official organizational communications created to support the critical principles of data privacy and data protection. Key policies include information security policy, privacy policy, and acceptable use policy. Each organization must evaluate its use of the Internet and its internal systems to determine the type of privacy policy it needs to protect all involved parties.

An organization practices *due care* by taking responsibility for all activities that take place in corporate facilities. An organization practices *due diligence* by consistently implementing and maintaining these security procedures to protect its facilities, assets, and workers. *Due process* guarantees that, in the event of a security issue by a worker, the worker receives an impartial and fair inquiry into the incident to ensure that the worker's employment rights are not violated.

Organizations employ several means for protecting information and information systems, including firewalls, intrusion prevention systems (IPS), security information and event management systems (SIEM), security orchestration and response (SOAR) systems, threat intelligence platforms (TIP), file integrity monitoring (FIM), file activity monitoring (FAM), and encryption. A data classification policy defines data sensitivity levels and handling and disposal procedures to ensure adequate protection of sensitive information.

Security incident response procedures (IRPs) ensure that organizations are equipped to monitor and detect events and incidents, and to respond to them quickly and effectively, thereby reducing their impact and scope.

Access controls ensure that only authorized personnel can read, write to, modify, add to, or delete data. They also ensure that only authorized personnel can access the different information systems and equipment used to store, process, transmit, and receive sensitive data. Effective physical access controls ensure that only authorized personnel are permitted to access work centers, processing centers, and other work areas, contributing to a better environment for information security and workplace safety. An organization's HR department is an important link regarding organizational and worker security. The HR department is responsible for hiring employees, tracking contractors, consultants, and other temporary workers, ensuring that workers comply with company codes and policies during their term of employment, and maintaining organizational security in the event of worker termination.

Information security awareness training ensures that an organization's workforce understands policies, behavioral expectations, and safe computer and Internet use. Organizations must determine the frequency, content, and audience for general and specialized training. Records substantiate that training has taken place.

Business continuity management is the overall effort to ensure that an organization's critical processes and systems continue operating despite natural or human-caused incidents and disasters. Resilience engineering ensures that information systems and processes maintain the required availability, even in the face of disruptive events.

Exam Essentials

Understand that information security supports risk management, not just technology protection. In CRISC, information security is evaluated based on how it reduces business risk, supports governance, and aligns with risk appetite and tolerance, rather than on technical implementation details alone.

Recognize that the primary goals of information security are confidentiality, integrity, and availability (CIA). Most security controls exist to directly or indirectly support one or more CIA objectives. Confidentiality ensures that only authorized personnel can access data.

Integrity ensures that only authorized changes may be made to data. Availability ensures that data and systems can be accessed on demand. Life safety is an emerging concern, but is not a formal part of the CIA triad.

Remember that security frameworks provide structure, not prescriptive solutions. Frameworks such as NIST RMF, NIST CSF 2.0, COBIT 2019, ISO/IEC 27001, and Risk IT serve as models for organizations to build their own governance and management systems. These frameworks operate at the business level and are technology-agnostic.

Know that NIST RMF is a system-centric, prescriptive risk management framework. The NIST Risk Management Framework follows a seven-step lifecycle (Prepare, Categorize, Select, Implement, Assess, Authorize, Monitor) and is commonly used to manage risk across the system lifecycle.

Understand that NIST CSF emphasizes governance and enterprise risk alignment. The NIST Cybersecurity Framework (CSF) 2.0 introduces the Govern function and focuses on integrating cybersecurity risk with ERM, making it an outcome-oriented standard rather than a framework of controls.

Learn that COBIT supports governance and management of enterprise information and technology. COBIT is not a risk framework itself, but an IT process and governance framework that integrates risk management across business and IT processes.

Know that the Risk IT Framework focuses explicitly on IT-related risk. Developed by ISACA, Risk IT addresses risk governance, evaluation, and response, and complements COBIT by providing more granular risk management processes.

Know the distinct roles of policies, standards, guidelines, and practices. Policies define required outcomes and behavior, standards specify mandatory implementation details, guidelines provide guidance on policy implementation, and practices represent generally accepted methods that may or may not be mandatory.

Understand that access management is a critical control for confidentiality and integrity. Principles such as least privilege, separation of duties, access provisioning and deprovisioning, access reviews, identity proofing, and privileged access management are frequently tested concepts.

Remember that incident response is a corrective risk response capability. Incident response limits the impact of realized risk events through preparation, detection, containment, recovery, and post-incident review, and must align with governance and legal requirements.

Understand that security awareness and risk awareness training reduce human-related risk. Effective programs support risk culture, reinforce policy compliance, and provide measurable indicators that contribute to risk monitoring, reporting, and accountability.

Know that privacy creates distinct regulatory and business risk scenarios. Laws such as GDPR, CCPA/CPRA, and sectoral regulations impose mandatory controls, reporting obligations, and penalties, making privacy a key driver of governance, incident response, and third-party risk management.

Review Questions

The following questions are designed to test your understanding of this chapter's material. For more information on how to obtain additional questions, see this book's introduction. Answers can be found at the end of this chapter.

1. What is the primary objective of information security from a CRISC perspective?
   - **A.** Preventing all security incidents
   - **B.** Ensuring compliance with technical standards
   - **C.** Reducing risk to acceptable levels aligned with business objectives
   - **D.** Maximizing investment in security controls
2. What best distinguishes a security policy from a security standard?
   - **A.** Policies are optional; standards are mandatory
   - **B.** Policies define intent; standards define implementation requirements
   - **C.** Policies apply externally; standards apply internally
   - **D.** Policies are technical; standards are managerial
3. Which activity best supports effective access management governance?
   - **A.** Implementing multifactor authentication
   - **B.** Performing periodic access reviews
   - **C.** Encrypting authentication databases
   - **D.** Centralizing authentication services
4. Which principle most directly limits the risk associated with excessive collection of personal data?
   - **A.** Accuracy
   - **B.** Purpose limitation
   - **C.** Data minimization
   - **D.** Transparency
5. Which factor most increases the likelihood that a control will degrade over time?
   - **A.** Initial over-engineering
   - **B.** Stable business processes
   - **C.** Changes in technology and threat landscape
   - **D.** Strong policy enforcement

6. Which activity best supports ongoing control effectiveness?
   - **A.** One-time certification
   - **B.** Periodic testing and monitoring
   - **C.** Control documentation
   - **D.** Control automation

7. Which principle is NOT part of the CIA triad?
   - **A.** Accountability
   - **B.** Availability
   - **C.** Confidentiality
   - **D.** Integrity

8. Which role is primarily accountable for accepting residual risk?
   - **A.** Risk owner
   - **B.** Control owner
   - **C.** System administrator
   - **D.** Internal auditor

9. Which metric provides the strongest evidence that information security awareness training is influencing risk outcomes?
   - **A.** Number of training modules completed per employee
   - **B.** Trend in phishing simulation failure rates over time
   - **C.** Percentage of employees passing post-training quizzes
   - **D.** Cost of the training program per user

10. A privacy impact assessment identifies high risk, but no control changes are made. How should this be classified?
    - **A.** Risk avoidance
    - **B.** Risk transfer
    - **C.** Risk elimination
    - **D.** Risk acceptance

11. Which approach most directly aligns information security awareness training with enterprise risk management objectives?
    - **A.** Delivering standardized training content to all employees
    - **B.** Updating training materials annually based on industry trends
    - **C.** Tailoring training content to roles, threat scenarios, and risk exposure
    - **D.** Outsourcing training delivery to a third-party provider

12. Which practice best supports a risk-aware security culture?

- **A.** Integrating security awareness with risk awareness
- **B.** Increasing security tool deployment
- **C.** Limiting training to IT staff
- **D.** Publishing technical configuration guides

13. What distinguishes GDPR from many U.S. sectoral privacy laws?

- **A.** GDPR applies only to government entities
- **B.** GDPR is principles-based and extraterritorial
- **C.** GDPR excludes enforcement penalties
- **D.** GDPR focuses solely on breach notification

14. Which NIST CSF 2.0 Core Function was newly introduced?

- **A.** Identify
- **B.** Protect
- **C.** Govern
- **D.** Detect

15. Which scenario best demonstrates ineffective incident response governance?

- **A.** An untested incident response plan with unclear escalation authority
- **B.** An incident response plan integrated with BCM
- **C.** Regular tabletop exercises
- **D.** Automated alerting tools

16. Which outcome best demonstrates that an information security awareness training program is effective?

- **A.** Measurable improvement in secure user behavior aligned to risk scenarios
- **B.** Completion of mandatory annual training by all employees
- **C.** Reduction in the number of reported security incidents
- **D.** Increased enforcement of disciplinary actions for policy violations

17. Why does CRISC emphasize governance over technical detail in security frameworks?

- **A.** Governance replaces security controls
- **B.** Governance determines accountability, priorities, and risk decisions
- **C.** Technical controls are outside CRISC scope
- **D.** Governance ensures compliance automation

18. Which document typically specifies mandatory implementation details for a policy?

A. Standard

B. Guideline

C. Procedure

D. Framework

19. Which characteristic most increases privacy-related risk exposure?

A. High system availability

B. Minimal data retention

C. Excessive collection of personal data

D. Strong authentication controls

20. What is the primary risk addressed by incident response practices?

A. Impact of realized risk events

B. Threat modeling accuracy

C. Control design completeness

D. Vulnerability creation

21. Which document primarily defines how an organization will continue critical operations during a disruptive event?

A. Incident response plan

B. Business continuity plan

C. Risk register

D. Disaster recovery policy

22. Which activity is most critical for determining recovery priorities within a business continuity program?

A. Threat intelligence analysis

B. Vulnerability scanning

C. Control self-assessment

D. Business impact analysis (BIA)

23. Which metric is used to define the maximum acceptable amount of data loss following a disruption?

A. Recovery Point Objective (RPO)

B. Recovery Time Objective (RTO)

C. Mean Time to Recover (MTTR)

D. Service-level Objective (SLO)

24. Which scenario best demonstrates alignment between business continuity management and risk management?
    - **A.** Conducting annual penetration tests
    - **B.** Updating access control lists quarterly
    - **C.** Designing recovery strategies based on risk scenarios and business impact
    - **D.** Deploying endpoint detection and response tools

25. An organization meets its technical recovery objectives during a disaster but still experiences significant financial and reputational damage. What is the most likely root cause?
    - **A.** Inadequate disaster recovery testing
    - **B.** Failure to align BCM objectives with business expectations
    - **C.** Insufficient backup frequency
    - **D.** Lack of incident response automation

Answers to Review Questions

1. C. CRISC emphasizes risk reduction aligned with business objectives, achieved through risk treatment. CRISC does not directly address absolute prevention or technology optimization. Not all security incidents can be prevented. Compliance with technical standards and maximizing control investment are important, but not the primary CRISC drivers.

2. B. Policies state *what* must be achieved, while standards define *how* policies are implemented. Both policies and standards are mandatory and are applied according to the scope of the program.

3. B. Access reviews validate the ongoing appropriateness of access, and support governance and risk oversight. MFA, encryption, and central authentication services may reduce risk and improve compliance, but are not governance drivers.

4. C. Data minimization requires organizations to collect and retain only the personal data necessary for a defined purpose, reducing exposure and impact in the event of misuse or compromise. Purpose limitation addresses the ways in which PII is used.

5. C. Environmental change is a major cause of control degradation. Controls designed today to counter today's threats may be ineffective in the future when threats evolve.

6. B. Effectiveness depends on ongoing validation, not documentation alone. One-time certification is not enough: controls must be periodically examined to ensure they remain effective. Documentation is critical, while automation may make controls more efficient (and possibly more effective).

7. A. Accountability is important for governance, but is not part of the CIA triad. The CIA triad is confidentiality, integrity, and availability.

8. A. Risk owners are accountable for risk acceptance decisions. Control owners are responsible for operating controls established as a result of risk assessments. System administrators are not business decision-makers, but stewards of information systems. Internal auditors are independent parties that examine controls for effectiveness.

9. B. Changes in simulated phishing failure rates reflect real-world behavior and provide outcome-oriented evidence of reduced human-related security risk.

10. D. Identified risk without mitigation reflects risk acceptance, whether explicit or implicit. The other choices all involve some sort of operational change.

11. C. Risk-aligned security awareness training focuses on role-specific threats and exposure, ensuring that awareness efforts address the organization's most significant risk scenarios.

12. A. CRISC emphasizes behavioral and cultural risk reduction. The other choices may benefit the organization, but have little or no impact on culture.

13. B. GDPR applies broadly, even outside the EU, when EU data subjects are involved. GDPR applies to all parties, including government and industry, and includes enforcement penalties.

14. C. NIST CSF 2.0 added the Govern function to emphasize governance and ERM alignment. The other CSF core functions: Identify, Protect, Detect, Respond, and Recover are in both 1.x and 2.0 versions of the CSF.

15. A. Lack of testing and unclear authority undermine effective incident response and risk management. The other choices do not indicate ineffective governance.

16. A. The effectiveness of an information security awareness training program is demonstrated by behavior change that reduces risk, not merely by training completion or enforcement activity.

17. B. CRISC focuses on decision-making, ownership, and oversight rather than implementation details. Governance does not replace security controls; rather, it manages their introduction, operation, and effectiveness testing. Technical controls are enacted based on risk assessment results and risk treatment decisions. Compliance automation, while potentially helpful, is not relevant for CRISC.

18. A. Standards define mandatory requirements that support policy implementation. Guidelines represent optional implementation and/or operational guidance. While mandatory, procedures don't specify implementation details; they specify only operational activities. Frameworks are collections of policies, standards, or controls.

19. C. Collecting more personal data than necessary increases regulatory and breach impact risk. High availability, minimizing retention, and strong authentication are likely to reduce privacy-related exposure.

20. A. Incident response limits the impact of events after risk has materialized. Threat modeling and control design influence risk independent of incident response practices. Vulnerabilities exist and are discovered, regardless of any organization's incident response capabilities.

21. B. A business continuity plan (BCP) documents strategies and procedures to maintain or quickly resume critical business operations during and after disruptions. Any critical business operations requiring IT system support will also include disaster recovery plans for each such IT system.

22. D. A business impact analysis (BIA) identifies critical processes, acceptable downtime, dependencies, and the impact of disruptions, forming the basis for recovery priorities and resource allocation.

23. A. Recovery point objective (RPO) defines the point in time to which data must be restored, representing the maximum tolerable data loss measured in time.

24. C. Business continuity management aligns with risk management when recovery strategies are selected based on evaluated risk scenarios and their potential business impacts.

25. A. Technical recovery success alone is insufficient if BCM objectives (such as RTOs and RPOs) are not aligned with actual business tolerance for disruption, leading to unacceptable outcomes despite system restoration.

Appendix

Implementing and Managing a Risk Management Program



This appendix takes a more pragmatic view of building and managing a risk management program. Rather than the somewhat academic approach in the CRISC job practice, I draw directly on my experience building several risk management programs throughout my career.

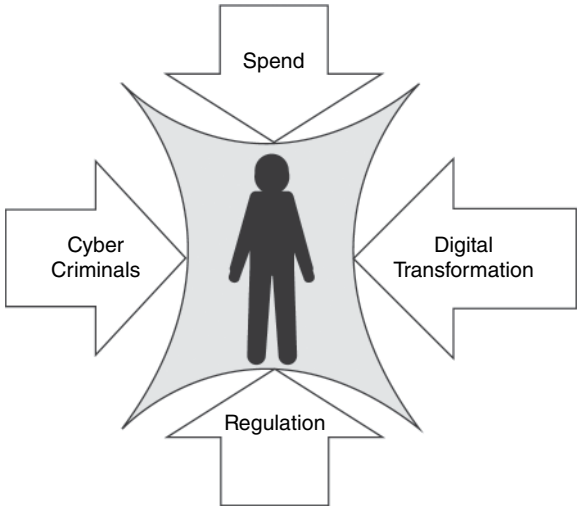
I have built or improved numerous risk management processes and programs. The method and technique are not overly complicated. The most challenging part of building a risk management process is figuring out how to engage with executives in the organization—so that they understand what the process is all about and participate meaningfully. After all, it's their business, and I just want to make sure they make informed decisions related to various types of risk.

Today's Risk Landscape

It's a crazy world for an IT or cyber risk leader. While the profession is both challenging and rewarding, increasing pressures on all sides are making the job more difficult (see Figure A.1). Some of the factors at play include the following:

- **Innovation and chutzpah of cybercrime:** Cybercriminal organizations are thriving. Ransomware, extortion, social engineering, and DDoS are highly successful, drawing even more into this dark industry. Intruders develop new attack methods in shorter periods of time, indicating the presence of increased research and development. Cybercriminals are using AI to supplement their tooling, and they are succeeding.
- **Upward pressure on regulations:** The success of cybercriminals is putting pressure on governments to pass laws requiring organizations (and governments) to do more to protect sensitive information. Industry sector regulations like the Gramm-Leach-Bliley Act (GLBA) and the Health Insurance Portability and Accountability Act (HIPAA), along with emerging privacy regulations like the European General Data Protection

FIGURE A.1 Pressure comes from all sides on security and risk programs today.



Regulation (GDPR) and the California Consumer Privacy Act (CCPA—amended by the California Privacy Rights Act—CPRA), are building layers of regulation on already strained businesses. These and other regulations are discussed in Chapter 5.

- **Downward pressure on spend:** Management is trying to keep a lid on rising costs and hold the line on spending increases while expecting security organizations to better defend against attacks.
- **Limited talent pool:** It's becoming harder to find qualified security and risk professionals to guide organizations out of the mess they're in. Organizations take months, half a year, or more to find new hires. I won't argue or debate the finer points on this matter, except to say that companies have a hard time finding the people they want to hire for their security teams.
- **Hybrid workforce:** Organizations are relying more on contractors, temps, and external service providers for projects and internal operations. As remote work grows, organizations are grappling with a broader range of workers across more locations and jurisdictions.
- **AI:** Proving itself to be the next industrial revolution, AI is upending industries from software to hospitality and beyond. Organizations are flocking to AI as they did to the cloud and the Internet, often without fully understanding the risks and consequences. They often fail to implement effective data governance and AI governance structures, resulting in additional unforeseen risks, including shadow AI, data leakage, hallucinations, privacy issues, and AI's opaqueness. Already, we've seen mega-corporations lay off thousands, only to hire them back again when their AI strategies failed.[1]
- **Increasing complexity of digital transformation:** Organizations' IT environments have never been more complex. Despite pressure to standardize on tools and technologies, IT organizations have more technologies, services, and operations to manage.
- **Introduction of IoT devices:** Increasingly, IT departments deal with more than just servers, desktops, laptops, printers, and mobile devices on their networks. Many other devices are plugged into internal networks, including "smart" devices, appliances, and voice-activated assistants. The so-called "attack surface" expands faster than a security team can ever hope to monitor and protect it.
- **BYOD:** Increasingly, organizations' workers use personally owned devices, mainly smartphones, to access and manage business information and systems. Such organizations have effectively ceded control of their information and information systems if these devices are not actively managed.
- **Cloud-based services:** These services reduce the effort and friction of acquiring new information systems. Sometimes, however, this goes too far, leading individuals and workgroups to unilaterally decide to employ cloud-based services, evading the scrutiny of legal and cybersecurity professionals. Corporate information has escaped the control of IT and risk owners in many cases. Also, organizations sometimes fail to implement

[1] For security and risk professionals with AI governance responsibilities, refer to my book, *AIGP Artificial Intelligence Governance Professional Study Guide* (Wiley Publishing).

processes and controls to manage cloud security in line with the shared responsibility models published by many cloud providers.

- **Increasing variety of security tooling:** Many thousands of security software companies compete in the marketplace with their tools. Security leaders are increasingly distracted by these vendors and their promises of stopping whatever the latest attack techniques happen to be at the moment.

These factors weigh heavily on security organizations, resulting in sub-par security monitoring and response, with too few personnel working on important issues.

Allegory of the Battlefield

Consider the photograph in Figure A.2 that shows Major General Alexander M. Weyand poring over a battlefield map in Korea.

Imagine, if you will, that General Weyand is taking in information from scouts on the enemy's position and discussing where and when to deploy troops to counter the enemy's moves and even gain territory. This is all about the allocation of scarce resources to achieve the maximum margin of victory.

The soldiers on the front line depend on this strategizing—their very lives depend on it. Political leaders rely on this strategizing as well, so that their armies will meet their objectives of countering the enemy's objectives.

FIGURE A.2 Major General Alexander M. Weyand poring over a battlefield map in Korea.



Source: Courtesy of U.S. National Archive.

Consider, for a moment, what would happen if an army lacked this planning and decision-making function. Without it, lower-level officers would be commanding their regiments to go this way and that, based on what they could see within their limited purview. Lacking big-picture information and purpose, they may succeed or fail, but they won't be operating with a clear, central purpose.

Risk management is the modern equivalent of those field commanders. In risk management, security specialists are examining threat intel, considering available internal capabilities, and deciding which of many issues are the most important to discuss and act on. Formal risk management is the modern equivalent of field commanders studying the enemy's plans, available resources, and most effectively applying their resources to meet their objectives.

Most organizations do not do this, at least not in any repeatable way. Instead, when decisions are made, they usually lack any analysis or triage to determine the best way forward on any given issue. The result: high-risk matters get little attention, while other issues get more than their fair share. Scarce resources are applied to the wrong issues, while the highest-risk issues are poorly resolved or remain entirely unresolved.

What Is a Risk Management Program?

To fully describe a risk management program, you first need to know what a *program* is. In business today, a program is a set of defined roles and responsibilities, as well as processes or procedures that are carried out indefinitely. A program consists of multiple supporting projects supporting a defined effort.

A risk management program is simply a process or set of processes through which

- Risks are discovered
- Risks are cataloged
- Risks are analyzed
- Business decisions about how to address risks are made

Also, follow-up is performed to ensure these business decisions are carried out on time and as agreed. The primary artifacts found in a typical risk management program include the following:

- **Program charter:** This document defines the purpose and scope of the risk management program, its roles and responsibilities, interfaces with other programs and processes, and its processes and records.
- **Process or procedures:** This document describes the steps in various risk management processes and procedures, who carries them out, and what records are maintained.
- **Risk register:** The central repository of identified risks. The risk register is one of the historical records of a risk management program.
- **Risk analyses:** Individual risks are often studied in greater detail, with additional analyses to help decision-makers decide what to do about these risks.
- **Business records:** Meeting minutes, risk treatment decisions, metrics, reports, and other actions taken.

So far, this may sound formal. Risk management should be consistent and methodical so that its proceedings are managed and recorded.

The Purpose of a Risk Management Program

Many kinds of threats and risks demand the attention of cybersecurity leaders. Even the best-equipped organizations have insufficient resources to deal effectively with every identified risk. A risk management program facilitates informed decision-making, helping to prioritize resources to deal with identified risks.

When properly designed and implemented, a risk management program serves as a *triage* apparatus to help identify and distinguish more significant risks from lesser risks. Instead of going into the details of determining the actual probability and financial impact of threats, risk management's triage sorts out risks to help management understand which risks are more critical to address sooner and which can be addressed later.

A common mistake in risk management programs is the time spent on detailed analysis to determine accurate probabilities and actual costs of threat events. While it may be helpful at times to estimate probability and impact, this is not the sole purpose of risk management. Instead, the triage function of risk management identifies risks needing prompt attention versus those whose attention can wait.

In this respect, a qualitative analysis method can be far more useful than a quantitative one. If risks are still considered "very high," it really doesn't matter whether it's 85.2% versus 95.7%; it's still coming, and it's still bad. And if it's the biggest risk, it's the biggest risk, regardless of the precision of its probability or impact.

When properly designed and implemented, risk management identifies the most critical risks and facilitates business-level decision-making by business leaders to address them.

Increasingly, the absence of an effective risk management program may represent negligence and a failure of due care. Few business executives tolerate their organizations being in such a state.



Not Only IT Risk Management

Many in the business think of risk management as consisting primarily of cyber risk and increasingly of privacy risk. The scope of risk in the CRISC certification consists of all IT risks, which goes beyond cyber to include various aspects of operational risk, such as the capacity and availability of IT systems to continue running properly.

The methodology of risk management, however, transcends IT altogether. Other risk domains include workforce risk, economic risk, market/competitive risk, currency exchange risk, geopolitical risk, natural hazards risk, fraud risk, and more. The techniques used in risk management are agnostic to the nature of risk.

Risk management is risk management, regardless of the focus or the context. It's only necessary for a risk management program to have someone skilled in the techniques of risk management and someone skilled in the subject matter of the risk management program. The program can succeed, whether as a single person or two people.

Risk Management Program Foundations

Risk management decisions are often consequential and may even alter the trajectory of an organization. Because risk management is so important, you should consider developing some formal artifacts that define the program and some of its operating parameters, including:

- **Risk management policy:** A formal policy document that describes the program, including required records and activities, such as:
  - Risk assessments: How often they must be performed, by whom, and using what methodologies or practices
  - Risk treatment parameters
  - Escalation and exception procedures, approvers, and records
  - Quarterly executive risk reviews and board reports
  - Business records and where they may be found
  - Roles and responsibilities, including the composition of the Three Lines of Defense, and the levels of approval required for various levels and types of risk
  - Interfaces with other risk management programs such as ERM
- **Risk appetite and risk tolerance:** A formal risk appetite and risk tolerance statement serves as guardrails to guide executives as they consider individual risk treatment decisions. While risk appetite and tolerance statements rarely rise to the level of "thou shalt not ever accept a risk that is greater than X," they serve as a proxy for decision precedents.

Like other formal documents, the ones listed here should be managed formally, including periodic reviews (annually is suggested) and revisions, approved by appropriate levels of management.

The Risk Management Lifecycle

Like many business activities, risk management is a lifecycle process. Risks are identified in various ways, and specific activities are performed for each risk. The entire lifecycle is described in detail in this section.

Risk Discovery

Chapter 2 lists how new IT risks can be discovered, with complete explanations for each. I provide just the summary here:

- Audits
- Penetration tests
- Security advisories

- Whistleblowers
- Threat modeling
- Risk assessments
- Security and privacy incidents
- Operational incidents
- News and social media articles
- Law enforcement
- Regulators
- Professional networking
- Passive observation
- Risk-aware culture

This list should not be construed as complete. Instead, it should help you understand that there are many ways in which a risk manager will discover new risks. Put another way, it would be dangerous to limit the identification of new risks to a fixed set of inputs, as this could artificially eliminate a source that is critical for some. Risk discovery is sometimes methodical and sometimes spontaneous. Because of the spontaneity of risk discovery, you should consider all possible likely and unlikely sources of risk. I can no more tell you which risk sources to include than I can tell you which of those sources will be related to your next risk event.



My frequent use of the term *risk manager* signifies a *role* rather than a job title. The job title of risk manager is generally seen only in larger organizations, whereas those with other job titles fulfill the role.

The Risk Register

The *risk register* (sometimes known as a *risk ledger*) is the centerpiece of business records in a risk management program. Each entry in the risk register represents a different IT risk that has been identified at some point in time, as recently as earlier today or as distant as a year or two in the past.

Since this appendix focuses on establishing a risk management program, I recommend implementing an initial risk register using a spreadsheet program such as Microsoft Excel or Google Sheets. It's my style to build a program or process using largely manual procedures, then introduce automation later.

A risk register starts quite simply, with columns like these:

- Risk name
- Risk description in terms of the risk context
- Threat description
- Vulnerability description

As you accumulate records in the risk register, you'll find it helpful to add more columns. These columns indicate when the entry was recorded, its context, and a business-centric explanation:

- Entry date
- Entered by
- Risk description in business terms
- Risk context

While the description of risk in business terms appears to be just another bullet item, it is perhaps the most important field in the risk register. Without it, risk managers will have a more difficult time describing the risk in business terms to executives who may be asked to decide its disposition. Including this column in the risk register compels the risk manager to begin thinking about the risk in business terms early, giving them time to ponder and develop the risk story.

Rating risks for severity is a topic all its own, discussed in more detail in the following subsections.

You will need to add more columns as you move through the lifecycle and make risk treatment decisions. These columns describe basic risk treatment parameters:

- Risk status
- Risk treatment selection
- Risk treatment description
- Risk treatment owner
- Risk treatment commit date

As risk treatment activities reach their conclusions, you'll need additional columns:

- Risk treatment completion date
- Risk closure date

Over time, you'll be spending considerable time in the risk register. You'll want to view risks in various contexts and categories at times. Some columns that help classify risks include the following:

- Location
- Business unit
- Service provider
- Technology type

The technology type may consist of one or more columns that enable you to view risks in discrete categories. For instance, you may wish to view all risks related to mobile devices, endpoints, specific types of servers, database management systems, network devices, or other components. A service provider column can be used to label risks related to external parties.



Risk managers should resist the temptation to be overly granular in their risk register, or else they may find themselves spending too much time on the details and not enough time on risk analysis and risk treatment activities.

Rating Risks—Qualitative

Not all risks are created equal. Some may be nearly trivial, while others represent an existential threat to the organization. A risk register must contain some columns used to describe the level of risk for each entry.

The two basic qualitative measures of risk are probability and impact, as detailed next:

- **Probability:** This specifies the likelihood of the risk actually occurring. In qualitative terms, probability is typically expressed as High, Medium, or Low, or in basic numeric terms such as 1 (Low) to 3 (High), 1 to 5, or 1 to 10.
- **Impact:** This specifies the impact on the organization due to an actual occurrence of the risk. Like probability, impact in qualitative risk ratings is expressed as High, Medium, or Low or on a Likert numeric scale such as 1–3, 1–5, or 1–10.

Note that, in qualitative risk rating, no attempt at an actual probability or impact is being made. Instead, the purpose of qualitative risk ratings is *triage*—these ratings help the risk manager distinguish entries with higher risks from entries with lower risks.



Qualitative Terms Help with Ratings

Rating the probability and impact of risks on Likert scales of 1–3 or 1–5 can, at times, be too abstract. For this reason, I recommend a list of terms be associated with these levels for both probability and impact. Using these terms as a guide can help the risk manager perform more consistent risk ratings for risk register entries.

For event probability on a scale of 1–5, these terms can be used:

- 5 = Certain
- 4 = Likely
- 3 = Possible
- 2 = Unlikely
- 1 = Rare

Similarly, these terms can be used for impact:

- 5 = Global
- 4 = Country
- 3 = Department
- 2 = Team
- 1 = Individual

These impact terms describe the scope of impact for a risk event. The highest-rated events would affect the entire organization, while the lowest-rated events would affect only a single worker.

Two additional qualitative measures can be introduced that help the risk manager better understand the overall risk for each entry:

- **Asset value:** This qualitative term expresses the value of the asset(s) involved in a specific risk. Like other qualitative terms, asset value is expressed as High, Medium, or Low, or on a simple numeric scale. Generally, this represents an asset that would have to be recovered or replaced should a relevant threat occur.
- **Operational criticality:** This term expresses the operational importance of an asset or function. An asset that the organization depends on for high-value operations (even if the value of the asset itself is low) would be assigned a higher value.

While these two measures are not seen as often in qualitative risk analysis, they help distinguish risks in that all-important triage to identify the most critical risks.

Calculating Overall Risk

The four qualitative risks described here can help the risk manager categorize risks in the risk register. One final column can help tie them all together, which would simply be called "Risk." This column can be easily calculated by multiplying the ratings of the four risk columns together (you could reasonably choose to *add* the values instead of multiplying them). For example, a risk item in the risk register is rated as follows:

- Probability = 4
- Impact = 2
- Asset value = 2
- Operational criticality = 1

The risk value in this example is $4 \times 2 \times 2 \times 1 = 16$.



Beware of the trap of false precision in qualitative and semi-quantitative risk analysis. Remember that the purpose is *triage,* not absolute risk values.

Using Qualitative Risk Ratings

In qualitative risk analysis, the individual risk columns, as well as the overall risk column, help the risk manager better understand how various risks compare. No attempt is made to calculate the actual probability of occurrence or the costs incurred if a risk event occurs. Instead, qualitative risk ratings help the risk manager understand which risks warrant more focus and attention.

Other columns in the risk register provide opportunities for focus, including those columns specifying the type of asset or system, the business unit, the geographic location, and more. The ability to focus on a portion of the risk register associated with a specific part of the organization, a particular location, or other distinguishing characteristics helps the risk manager focus on which risks matter most versus which matter less.

In a sample risk management scenario, a risk manager is preparing to have a conversation with the part of IT that manages end-user endpoints—primarily laptop and desktop computers—in a particular region of the world. To prepare, the risk manager can filter the risk register to view

just the risks associated with end-user endpoints, as well as risks in particular global regions. The resulting risk will contain just those items, which the risk manager can then examine more deeply to distinguish the items of greater concern from those of lesser concern. Now, the risk manager can have a more meaningful conversation with the IT leaders about risks directly associated with their work.

Rating Risks—Quantitative

Sometimes, a risk manager is asked to go beyond qualitative risk rating and provide actual probabilities and financial impacts of risks. "How much will it cost us if customer data is stolen?" is a reasonable question that an executive may pose to a risk manager. Arguably, the risk manager should provide an estimate to help the executive better understand the consequences of a security breach.

In my practice, I do not attempt to uplift every item in the risk register from qualitative to quantitative terms. Doing so would take a considerable amount of time and not result in much additional insight. Instead, it may be better to embark on quantitative risk analysis when examining individual risks. I discuss this further in the section, "Performing Deeper Risk Analysis."



The Risk Register Is a Living Document

The risk register is a living document, and even when risks are reduced, this is recorded and reflected in the document; they are not simply deleted. The risk register provides a history of risk management within the organization. The contents of the risk register should be managed accordingly.

Types of Risk Registers

Many kinds of risks are present everywhere in an organization, and it makes little sense to put them all together in one risk register. Imagine if there was a risk entry associated with currency exchange risk, and adjacent to it is another entry that expresses the misconfiguration of a server. Both are useful, even valuable, but they differ greatly in their context and importance. Further, particularly in larger organizations, other individuals will manage these risks and discuss them with different executives.

Layers of Risk Registers

Different levels of risk are typically maintained in separate risk registers. Most often, the following three risk registers will exist:

- **Entity-level risk:** These are the overall "big picture" risks that affect an entire organization. The types of subject matter in an entity-level risk register generally include the following:
  - Workforce-related risks
  - Currency exchange risks in organizations operating in multiple countries, or with suppliers in other countries
  - Economic risks having to do with macroeconomic business cycles
  - Risks that exist as an aspect of organizational culture

- **Program-level risk:** These are risks most often associated with business processes throughout the organization. Examples of the types of risks found at this level include the following:
  - **Effectiveness of vulnerability management:** If, for instance, the IT department is lackadaisical in its security patching, an entry may be created in the risk register that describes this systemic problem.
  - **Identity and access management:** An organization may have little discipline in its access request, fulfillment, and termination procedures.
  - **BYOD:** An organization may not have controls to prevent employees' use of personally owned devices for managing critical data, leading to a significant data leakage risk.
- **Asset-level risk:** These risks are associated with individual assets and groups of assets. Often, these risks are identified with scanning tools and other tooling. Examples of risks at this level include the following:
  - Misconfiguration of critical devices
  - Missing critical patches
  - Newly discovered unmanaged devices

Communicating Between the Layers

A single risk manager may be involved in two or all three of these risk registers in smaller organizations. More often, however, different persons manage each risk register separately. One might think that these risk registers would be isolated from each other. While this may be true, organizations are wise to be cognizant of opportunities for communication between those who manage various risk registers.

Classic examples of the crossover between risk registers include the following:

- Chronic occurrence of an IT department exceeding patching SLAs may compel a risk manager to create a program-level risk describing a risk with the operational process of patching.
- Chronic occurrence of a product engineering group avoiding engagement with security and privacy teams may compel a risk manager to create an entity-level risk describing a culture that lacks the essential cultural ingredient of "security and privacy by design."

The most frequent communication between risk register layers is a chronic or repeated failure of processes at lower layers that suggests a systemic problem at a higher layer.

Reviewing the Risk Register

Risk managers should have a thorough understanding of the contents of their risk register(s). This requires a risk manager to frequently review the risk register, using various filtering ("slicing and dicing") techniques to view subsets of the entire register. This attention could approach "preening," where a risk manager examines and periodically adjusts contents in the risk register.

Most often, a risk manager will occasionally adjust an aspect of risk rating of a risk register item. As the risk manager becomes intimately familiar with the risk register, they may realize that a particular risk entry's risk rating should be adjusted a bit up or down, making risk ratings more credible and reasonable.



A risk program aspiring to higher maturity should require recording rating changes and other changes to the risk register, or perhaps use a change control process requiring two or more persons to understand and agree to changes to the risk register.

In an example described earlier in this appendix, a risk manager is preparing for a conversation with a particular leader by examining risk register items relevant to that leader's mission. This is a helpful example, but a risk manager should not wait for a conversation to focus on those relevant risks. Instead, the risk manager should routinely examine smaller sets of risks in logical groupings such as location, type of asset, and business unit to become even more familiar with the risk landscape in different parts of the organization.

Formal Risk Register Reviews

While the risk manager may be spending quality time alone with the risk register, formal reviews with one or more additional people are warranted. These reviews can help improve the quality of the risk register through consensus. Reviews strengthen the integrity of the risk register, as they will represent the opinion of a group rather than a single individual.

Few other persons may have skills or experience in risk management in smaller organizations. Risk register reviews with others are still valuable: as the risk manager describes risks to others, opportunities for improvements can still be identified.

Better organizations will document risk register reviews and record any changes to risk register entries.

Reviewing Risks with Staff

It is often appropriate for the risk leader to review the risk register with other security, privacy, and risk staff. Generally, these discussions will consist of a closer look at the details in the risk register, including ratings and any completed risk analyses that include potential risk treatment options.

Staff members can be a good source of new material for the risk register. Reviews may, at times, remind staff members of other matters that may be good candidates for the risk register. Frontline personnel are often aware of good and poor practices in technology operations and business processes; earning their trust and using them as a source of intel on new risks can be a valuable pursuit.

When possible, staff reviews should not be limited to personnel in cybersecurity but could include staff from IT and other departments.



Risk-averse organizations may consider using reminders of confidentiality before reviewing risk register contents with staff.

Reviewing Risks with Business Leaders

At times, it is appropriate for the risk leader to review the risk register with business leaders. This discussion can take many forms, including the following:

- A detailed review of all relevant risks
- A high-level review of the "Top N" risks
- A high-level review of the nature of top risks

These reviews may be considered formal or informal, depending on many factors, including the maturity of the risk program and governance in the organization.

These reviews are not a part of risk treatment but instead are an assessment of the universe of relevant risks. These reviews help business leaders be better informed of the risks that have been identified, and possibly would include a discussion on what risks may be the subject of focus in the near term.

Identifying Candidates for Deeper Analysis

The purpose of risk register reviews is to reach an agreement on each risk's ratings and other data, as well as overall risk reduction in the organization. Risk reduction is accomplished by selecting risks that are considered suitable candidates for next steps: deeper analysis and, eventually, risk treatment.

The purpose of deeper analysis is to help the risk manager decide whether a specific risk is ready for risk treatment—the decision made by business leaders on the risk's final disposition. Before that decision can be made, the risk manager must first analyze the risk in more detail.

In the risk management program, the risk manager may establish formal criteria or rely more on gut instinct for selecting risks to study in more detail. A blended approach will likely be used in this selection, where risks meet criteria and are more "visible" than others. This visibility may be related to applicable regulations, internal politics, market reputation, service profitability, or other factors.

A risk manager may select less politically charged risks in a new risk management program. Business leaders may then be more comfortable making risk treatment decisions due to the absence of emotional or political factors. When business leaders are more familiar with the risk treatment process, risks with greater visibility (whatever that may mean) can be selected and treated.

Performing Deeper Risk Analysis

When a risk has been selected for deeper analysis, it's time for the risk manager to begin working on it. The purpose of this deeper analysis is to determine one or more viable risk treatment options that, hopefully, business owners will agree with, ratify, and support with adequate resources.

Individual risks are generally complex: discerning risk treatment options may require much thought and discussion. I prefer the ideal small-group discussion, where the risk manager and one or two subject-matter experts brainstorm to identify feasible risk treatment options.

There is no ideal number of options to analyze, although pragmatically, I recommend three to five options to be considered. If more than five options are selected, risk analysis will take additional time; with fewer than three, management may feel that too few options were considered.

In addition to cost, time, and effort, non-financial constraints often determine which risk treatment options are viable. Some options may be inexpensive, yet infeasible due to regulatory obligations, contractual commitments, reputational exposure, or ethical considerations. For example, a low-cost option that technically reduces risk may still be unacceptable if it conflicts with regulatory expectations, undermines customer trust, or creates adverse public perception if disclosed.

These constraints should be considered explicitly during analysis, even if they are not easily expressed numerically. Ignoring them can result in recommendations that appear rational on paper, but are rejected by business leaders or compliance functions. Over time, incorporating these considerations improves the credibility of risk analysis and reduces friction in risk treatment discussions.



In our continuing discussion of risk management, risk analysis, and risk treatment, I use terms that imply a simple spreadsheet program for recording risks and performing risk analysis. Using a spreadsheet may be practical for an emerging risk management program.

The Risk Analysis Worksheet

The worksheet used for additional analysis compares the untreated risk score with the risk scores for each of several possible risk treatment options. I suggest using a simple model, as described here.

Untreated risk:

- The same columns from the risk register, including likelihood, impact, criticality, asset value, and risk score.

Treated risk:

- An additional set of risk score columns, including likelihood, impact, criticality, asset value, and risk score.
- Additional columns that express the time, effort, and cost required to perform the risk treatment.
- Columns that express the reduction in risk and residual risk.

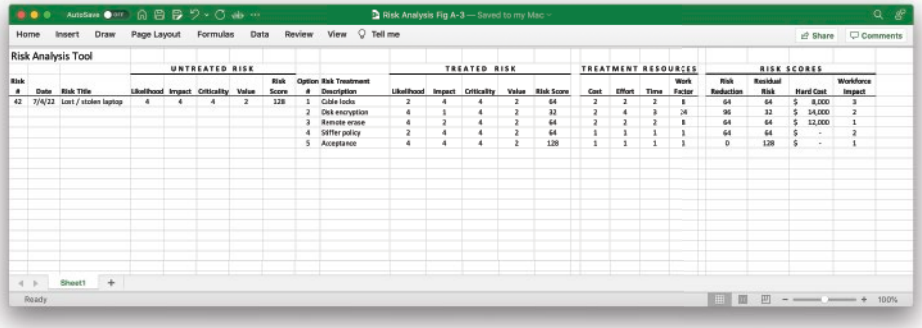
Risk treatment:

- One or more columns that would specify who or what department or team would perform the risk treatment.

Each risk treatment option occupies a new row in the worksheet. Each option is scored, with the scores representing the levels of likelihood, impact, criticality, and asset value, as well as the newly calculated risk score, as though the risk treatment were in place.

Once each risk treatment option is scored, the risk reduction and residual risk are shown. The risk reduction and residual risk, together with other columns showing hard costs and impact on the workforce, can help the risk manager better understand the feasibility of each risk treatment option.

FIGURE A.3 A simple risk analysis worksheet showing risk treatment options.



Depending on the risk's complexity, a detailed analysis may require many hours of research and study. The amount of time available for this work will depend on the risk manager's availability compared with other tasks. In some cases, detailed risk analysis may be considered a luxury.

Figure A.3 depicts a simple risk analysis worksheet that includes the initial risk rating, plus a description of various risk treatment options with their ratings.



I prefer to use a separate workbook for performing detailed risk analysis; otherwise, the central risk register workbook will soon be overcrowded with numerous risk analysis tabs.

Qualitative and Quantitative Risk

For risks that involve relatively low costs, it may be sufficient for the risk manager to rely only on qualitative risk analysis techniques to portray which treatment options are more effective than others at reducing risk. However, when the costs of the risk and risk treatment are high, the risk manager may need to perform a quantitative risk analysis to estimate actual risk treatment costs. Quantitative risk analysis is instrumental when the risk treatment option includes significant hard and/or soft costs.

Selecting Risk Treatment Options

For the risk analysis to contribute to successful risk treatment, the risk manager needs to understand how the business leaders in their organization think about risk. The risk manager should select one or more risk treatment options that business leaders can understand. Business leaders need to feel that their decision is not forced or coerced, and that the options represent a range of reasonable options.

In practice, the political feasibility of a risk treatment option often determines whether it will succeed, regardless of its technical merit. An option that requires significant behavioral change, reallocates budget from a powerful business unit, or implies prior management failure may encounter resistance that outweighs its risk-reduction value.

You will need to assess not only whether a risk treatment option should work, but also whether it can work in the organization's current political and cultural environment. Options that align with existing incentives, leadership priorities, and organizational momentum are more likely to be approved and executed. Conversely, options that challenge entrenched interests may require executive sponsorship or phased implementation to be viable.

As the risk manager ponders risk treatment options, they need to remember that the goal is to bring the elements together to tell a story.

Considering Risk Appetite and Risk Tolerance

Organizations are all over the place when it comes to risk appetite and risk tolerance. It is my experience (through many years of security leadership and security consulting) that most organizations have no formal statement or practice of risk appetite. In other words, management has not drawn a line in the sand stating that risks above a certain magnitude should not be routinely accepted.

Organizations lacking a formal risk appetite policy or statement must rely instead on gut feel for determining whether to accept any particular risk. While not ideal, as long as the risk manager properly facilitates the risk treatment conversation, business leaders can still make an informed risk decision, even without formal risk guardrails. In these organizations, risk appetite becomes apparent through the track record of risk treatment decisions.

A better situation is when the organization has formally defined risk appetite, which states the highest risks it is permitted to accept. I advise caution here: risk managers should resist the urge to reclassify specific risks to a lower level, just to facilitate acceptance. Suppose a particular risk exceeds the allowable risk appetite. In that case, the organization's *risk tolerance* comes into play: just how much more risk will the organization accept on an exception basis, and who is required to approve such an exception?

Understand that we're still in the risk analysis phase, but are anticipating possible outcomes for risk treatment discussions and an eventual decision. Risk managers should be prepared with talking points to argue for and against each risk treatment option, as well as the array of risk mitigation options. Remember that it is not the risk manager's responsibility to make risk decisions but rather to facilitate discussions leading to business leaders making informed decisions.

The Precedence of Prior Risk Treatment Decisions

In organizations without a well-defined risk appetite, documented precedent often becomes the de facto guide for future decisions. Past risk treatment decisions—particularly those approved by senior leadership—establish informal boundaries for what levels of risk are considered acceptable (or were at a point in time).

This underscores the importance of documenting risk decisions, alongside publishing policy. Recording what risks were accepted, mitigated, or rejected, under what circumstances, and by whom, creates a practical reference point for future deliberations. Over time, these precedents form an implicit risk appetite that business leaders recognize and rely upon, even if it is never formally articulated.

When precedent is undocumented, risk managers are forced to re-litigate similar risks repeatedly, increasing friction and inconsistency. When precedent is recorded and accessible, risk discussions become faster, more coherent, and more defensible.

Developing a Risk Treatment Recommendation

You should think about the risk analysis of a risk register item as the development of an entire package of information to be presented to senior management. The package will typically consist of the following:

- A brief description of the risk
- The context of the risk
- Typical risk scenarios, with their probability and impact, if and when realized
- Risk treatment options
- The recommended risk treatment option
- Backup information should management wish to explore the matter further

When presenting risk treatment options to senior management, you can reduce this to a single presentation slide and speak to the matter in business terms. Typically, there will be some discussion and a few questions. Depending on the organization's formality, there may be a formal vote or verbal support for or against a decision.

It's vital to develop talking points, anticipate questions, and rehearse the briefing with a trusted peer. You generally have one opportunity to tell the story of the risk, its potential impact on the organization, and what you propose be done about it.



It's always better to be prepared with more information than is used in a risk treatment proceeding.



Risk Treatment Overruled

As a risk manager in a public company, I brought to senior leadership in one of our risk management reviews a matter of workplace security and safety: a particular business office had no access control that segregated work areas from the reception area. Visitors during business hours could simply walk through the door and into the entire office suite, where about 100 employees worked. My recommendation was to accept this risk.

The company president stood up and said, "No, I will not accept this risk. We should have a wall and door constructed to keep visitors out of the work area unless they are permitted to enter."

While this may at first appear to be a correction or a reprimand, I took it as an engaged senior executive taking a different point of view and making an informed risk decision. At this point, I knew that senior leaders were engaged and not merely rubber-stamping my recommendations.

The Risk Treatment Decision

Up to this point, the risk manager has been analyzing one or more risks, developing viable risk treatment and risk mitigation options, and preparing a package to present to senior management. During the meeting in which the risk manager presents the risk and the recommended course of action (or inaction, in the case of risk acceptance), senior leaders should be allowed to ask questions before being asked to ratify the recommendation (or decide otherwise). Note the questions, who asked them, and how they were answered—this will help in future proceedings.

The call for the actual decision may range from formal to informal. An accurate accounting of the decision is critical, as it should be a part of officially published meeting minutes used to document the decision. Note those present and (optionally) who agreed and who disagreed. Risk management frameworks have cited four risk treatment options: mitigation, acceptance, transference, and avoidance.



In some cases, it may be wise to discuss the risk treatment matter with senior executives in one-on-one discussions, where they may be more apt to ask questions to understand the situation better. These discussions will also help the risk manager better prepare for the risk treatment discussion, and in some cases, it may even sway the risk manager to pursue a different remedy.

Risk Mitigation

When the risk treatment decision is *mitigation*, the decision should also include the naming of the person responsible for performing the risk treatment and executives' desire for the timing of risk treatment. This is particularly important if the risk treatment owner is not present in the discussion.

Any discussion about risk treatment options should be captured so that the risk treatment owner will have additional background and insight into why the particular risk treatment option was selected. This may help the risk treatment owner proceed more quickly with planning and execution without circling back to decision-makers to better understand the rationale behind the decision.

All the basic facts about risk mitigation (who made the decision, when the decision was made, who the risk treatment owner is, when the risk treatment should be completed, and so on) should be entered into the risk register. Then, the risk manager can conduct effective follow-up over the coming weeks, months, and quarters and report on the progress to the decision-makers.

Risk Acceptance

If the risk treatment decision is *acceptance*, it should not be perpetual but instead have an expiration date. After that expiration date, the risk will be reopened, and a new risk treatment decision will be made. The reason for expiring risk acceptance is this: conditions related to the risk may change in a year, so much so that a different decision may be needed in the future. Many things can change in a year:

- How the organization uses information systems
- The organization's reliance on critical business processes on information systems
- Changes in underlying information technology

- Changes in threats
- Changes in tools that help mitigate threats
- Changes in business priorities

One year may seem like a short period, but a lot can happen in a year. Make sure executives are on board with this concept.

Organizations with formal risk management programs sometimes utilize a *risk acceptance letter*—a signed memo in which a business unit leader or department head chooses to accept a specific, identified risk versus another risk treatment option, such as mitigate, transfer, or avoid. In the risk acceptance letter, the signatory accepts the potential consequences if an event related to the risk occurs in the future.



CISOs Are Not Risk Owners

Many organizations have implicitly adopted the mistaken notion that the chief information security officer (CISO) is the de facto risk owner for all identified risk matters.

In a properly run risk management program, risk owners are business unit leaders and department heads who own the business activity where a risk has been identified. For instance, if a risk is identified regarding the long-term storage of full credit card numbers in an e-commerce environment, the risk owner would be the executive who runs the e-commerce function. That executive would decide to mitigate, avoid, transfer, or accept that risk.

The role of the CISO is to operate the risk management program and facilitate discussions and risk treatment decisions, not to make those risk treatment decisions. A CISO can be considered a risk *facilitator* but not a risk owner.

Even when embraced and practiced, this concept does not always stop an organization from sacking the CISO should a breach occur. A dismissal might even be appropriate, for example, if the risk management program that the CISO operated was not performing as expected.

Risk Transference

One of the four basic risk treatment options, *risk transfer* means that some outside organization is willing to take on the risk. Often, this comes in the form of cyber insurance, but it could just as easily be a *business process outsourcing (BPO)* arrangement in which a service provider is willing to accept certain risks.

I provide a note of caution here. While an organization can outsource risk, no organization can outsource accountability. For example, a successful service business specializing in water damage mitigation outsources its entire IT and IT security functions to a large service provider, including information systems where customers register and request one-time or recurring services. The IT service provider suffers a breach, exposing its customers' personal information

and leading to numerous cases of identity theft and fraud. The service business attempts to explain that the external service provider is the party responsible for the breach. However, the customers are unhappy, blame the service provider, and threaten to organize a class-action lawsuit for the damages.

In this example, the service business cannot absolve itself because the IT outsourcer's lapse resulted in the breach. The service business decided to use this particular IT outsourcer and cannot absolve itself of responsibility because of an unfavorable outcome.

Cyber insurance is similar to this example. An organization with cyber insurance may suffer a breach, and the insurer may compensate the organization for specific losses. However, the organization will have lost some goodwill and may have a long-lasting tarnished reputation that can take years to undo.

It is imperative that senior executives fully grasp the complete picture of risk transfer.

Risk Avoidance

Risk avoidance is often the least understood risk treatment option. With risk avoidance, the activity that causes the risk is itself discontinued.

For example, a risk assessment reveals the use of an unsanctioned cloud service provider to process sensitive information. End users decided to use a free online service to translate sensitive documents, including contracts. In its end-user licensing agreement (EULA), the service provider claims ownership over any content that it translates. When this matter was brought to light in a risk treatment discussion, business executives directed the organization to immediately discontinue use of the service provider, thus avoiding the risk.



The Fifth Option in Risk Treatment

For decades, risk management frameworks have cited the same four risk treatment options: accept, mitigate, transfer, and avoid. There is, however, a fifth option that some organizations select: *ignore the risk*.

Ignoring a risk situation is a choice, although it is not considered a wise choice. Ignoring a risk means doing nothing about it, not even making a decision about it. It amounts to little more than pretending the risk does not exist. It's off the books. It is not even added to a risk register for consideration, but it represents a risk situation, nonetheless.

In some cases, such as for minimal risk items, this may be perfectly acceptable. A theft of a paperclip may be too small to warrant inclusion in a risk register. It would probably be wise to leave this off a risk register unless there is a specific reason to add it. In some cases, listing minimal risk is very critical because compliance requirements dictate that specific risks be considered in risk evaluations. Developing the right level of detail for a risk register requires experience, an understanding of an organization's culture, and a keen sense of balance.

Organizations without risk management programs may implicitly ignore all risks, or many of them at least. Organizations might also be practicing informal and maybe even reckless risk management—risk management

by gut feel. Without a systematic framework for identifying risks, many are likely to go undiscovered. This practice could also be considered as ignoring risks through the implicit refusal to identify them and treat them properly.

It is essential to distinguish between risks intentionally accepted and those simply ignored. Accepted risks are identified, understood, and deliberately approved by an appropriate decision-maker, often with conditions or a review date. Ignored risks lack this visibility and accountability, and they are far more difficult to defend if they later materialize.

Note that ignoring risk, particularly when governance requires that you manage it, is usually a violation of the principles of due diligence and due care. Many organizations can be legally charged with "willful negligence" if they have a duty to manage risk, and they simply don't.

Publishing and Reporting

Like any formal business program, risk leaders publish periodic reports and metrics from the risk management program. The risk program needs to inform business executives on the state of risk in the organization and proceedings in the program. Several aspects of risk reporting are discussed here.

Risk Register

The risk register itself is considered a detailed, operational business record. Reporting about the risk register generally consists of high-level depictions of risk instead of reporting the details. Challenges in reporting risk include the following:

- **Trends:** Changes in the size of the risk register, or the aggregate or average risk levels, may be misleading, particularly in newer risk programs that still identify risk.
- **Noise:** Without proper scoring or prioritization, the sheer length of the risk register (which may contain a long tail of low-rated risks) may distract executives from focusing on the highest risks that warrant action.
- **Completeness:** A risk register should not be considered a complete record of risk, as there may be numerous undiscovered and unidentified risks.

New Risks

Risk leaders may opt to disclose the number of, or even some details about, new risks identified in the reporting cycle. This reporting is a good indication that the risk program continues to keep its risk radar operating and open to risk discovery.

Risk Review

Reporting and metrics on risk reviews indicate engagement and involvement. Executives should expect risk leaders to periodically engage business leaders and department heads throughout the organization, and may want to know how often (and with whom) this occurs.

Risk Treatment

The rate at which risk treatment decisions are made, and trends in risk treatment options, show that the risk leader continues to work through the risk register.

Risk Mitigation Follow-up

When mitigation is the chosen risk treatment for a given risk, the risk leader needs to track that risk mitigation and report on whether mitigation is performed on schedule. Mitigation activities that are behind schedule can be highlighted so that executives can understand where their directives are falling short and which are being completed on time.

Risk Renewals

As discussed earlier in this appendix, risk acceptance should not be perpetual but should expire. Risk leaders may want to publish, as a leading indicator, the number of risks nearing renewal. Those risks will be redeliberated and new decisions made, whether to continue to accept the risks for another year or take a different approach.

Risk Prioritization

In risk reporting, leaders often care only about the top-priority risks, such as the top 5, 10, or 20 risks. These may not be precisely the top risks, according to the risk manager. However, these priorities articulate risk appetite and risk tolerance, as well as leadership's culture and risk focus. All too often, risk prioritization with managers boils down to nothing more than cost, unfortunately.

In practice, a Top-N risk list often reflects leadership attention as much as objective risk ranking. The risks that rise to the top are usually the most visible, politically salient, or tied to current business priorities, rather than those with the highest inherent risk. While this does not invalidate Top-N reporting, it is important to recognize it as a management focus tool, rather than a definitive statement of overall risk exposure.

You should avoid using reports and metrics that create unintended or counterproductive incentives. Metrics that reward a shrinking risk register, fewer reported risks, or rapid risk closure can discourage risk discovery and honest reporting. Effective risk metrics should reinforce transparency and engagement, not the appearance of reduced risk.

In most organizations, detailed information about risk, particularly vulnerabilities and safeguards, is treated as highly confidential. If it became public, that information might allow the organization to be attacked. Further, this information would reveal clues about the organization's risk culture. This is especially true in government and healthcare, in my experience.

Index

Note: Page numbers in *italics* and **bold** refers to figures and tables respectively.

A

acceptable use policy (AUP), 235–236
access deprovisioning, 244
access management, 241–246
access provisioning, 244
access review, 195, 245
accountability, 18, 88, 104, 243
after-action review (AAR), 116
alternate processing sites, 191
application programming interface (API), 179
architecture review and design review, 51
artificial intelligence (AI), 119–120
  benefits of, 87–88
  governance and human oversight, 88
  integration with frameworks and standards, 87
  limitations and risks of, 88
  models, 212
  in risk analysis and evaluation, 86–87
  in risk identification, 86
assets, 9–10
  asset data, collecting and organizing, 12–13
  asset data, sources of, 11–12
  classification, 69–70
  identification, 10–11, 69–70
  management, 10, 174
  valuation, 69–70
audit, 51
authentication, 242
authorization, 242–243

B

backup-based recovery, 190
bottom-up risk scenario development, 48
bow-tie analysis, 77, *78*
business alignment, 4
business continuity management, 250–254
business continuity planning (BCP), 184, *184*
business impact analysis (BIA), 78–83, *79*, **82**, 188, 251
business processes, 8–9
business process outsourcing (BPO), 111
business–technology alignment, 165–166

C

California Consumer Privacy Act (CCPA), 259–260
California Privacy Rights Act (CPRA), 259–260
Center for Internet Security (CIS), 128
certification, CRISC
  applying for, xxxiv–xxxv
  exam, xxviii–xxxiv
  experience requirements, xxvi–xxvii
  onsite testing center, xxix
  remote proctored testing, xxix–xxx
  retaining, xxxv–xxxviii
  revocation of, xxxviii
change management, 112, 135, 172
cloud and virtualization technologies, 209
  cloud service models, 210
  data protection and storage risk, 211
  identity and access in cloud environments, 211
  provider dependency and resilience, 211
  shared responsibility model, 210
  virtual networks and segmentation, 210–211
  workload protection and configuration management, 211

cloud infrastructure, 179. *see also specific types*
cloud shared responsibility models, 107
  continuous oversight, 110
  division of responsibilities, 108–109
  governance/contractual considerations, 110
  risk ownership and control selection, 109
  shared responsibility, 108
COBIT 2019, 3, 224, 230–232, **231–232**
code review, 51
code scan, 51
confidentiality, 225–226
configuration management, 112, 135, 173
configuration management databases (CMDBs), 173
continuing professional education (CPE), xxxv
  maintenance fees, xxxviii
  requirements, xxxvi
  sample records, **xxxvii**
  tracking and submitting, xxxvi–xxxvii
Continuous Integration/Continuous Delivery (CI/CD), 184
continuous third-party monitoring, 105–106
control deficiencies and gaps, 43–44
control deficiency analysis, 40
controls, 131–139, 236–237
  analysis, 133–134
  continuous monitoring and assurance, 137–139
  design, 131–132
  documenting and publishing, 134–135
  effectiveness, 41–43, *42*, *43*
  and frameworks, 124–129, **125**, **127**
  implementation, 134–136
  rationalization, 129–131
  recommendations, 44–45
  selection, 132–133
  testing and effectiveness evaluation, 136–139
  types, 121–124, **122**, *124*
COSO Enterprise Risk Management (ERM) Framework, 16
cost risk, 180
criticality analysis (CA), 81–82, **82**
cyber risk governance committee, 6
Cybersecurity Framework (CSF, NIST), 228–229

D

database management systems (DBMSs), 178
data discovery, 51
data lifecycle management
  activities in process, 197
  data classification, 195–196, 256
  data governance, 194–196
  data ownership and stewardship, 195
data privacy, 254–255
  regulatory foundations, 258–261
  as risk scenario, 257–258
data protection, 256–257
data sensitivity, 256
Deming Cycle, 3, *3*
DevOps/DevSecOps pipelines, 183
disaster declaration and activation procedures, 189
disaster recovery management, 184–186
  alternate processing sites, 191
  backup-based recovery, 190
  in business continuity, 188
  business impact analysis integration, 188
  disaster declaration and activation procedures, 189
  objectives and scope, 187
  plan documentation, 189
  planning activities, 187
  plan testing, 192–194
  prioritized recovery sequencing, 192
  rebuild-from-code and automation, 191
  recovery point objective (RPO), 187–188
  recovery time objective (RTO), 187–188
  redundant or high-availability architectures, 191
  replication and mirroring, 191
  risk assessment and scenario identification, 188

strategies, 190–192
strategy development, 190
testing metrics, 193–194
training and preparedness, 189
types of DR testing, 192–193
disaster recovery planning (DRP), 184, *184*
due care, 238
due diligence, 238
due process, 238
dynamic application security testing (DAST), 206

E

emerging risk
factors, 117–118
management of, 116–120
emerging technologies, 116–117, 207–215
AI models, 212
API-first architectures and API security, 213
artificial intelligence (AI), 119–120
cloud and virtualization, 209–211
containerization, 212
edge computing, 214
IT/OT convergence, 214
Kubernetes and orchestration, 213
microservices architectures, 213
modern architectures, 212–215
quantum computing, 212
risk factors, 207–208
security, 208
serverless and function-based computing, 214–215
system maintenance, 208
threats, 208
vulnerabilities, 209
zero-trust architecture (ZTA), 214
emerging threats, 118
emerging vulnerabilities, 118–119
end-user services, 172
enterprise architecture (EA), 164–169
enterprise risk management (ERM), 15–16
EU AI Act, 259
event-tree analysis, 75–76, *76*
Everything-as-a-Service, 179, 210
external ratings, 105–106
extraterritorial laws, 20

F

Factor Analysis of Information Risk (FAIR), 76–77
Family Educational Rights and Privacy Act (FERPA), 261
fault-tree analysis, 75–76, *76*
formal risk register reviews, 284

G

gap assessment, 50
General Data Protection Regulation (GDPR), 258–259
governance, 1–2
and human oversight, 88
organizational governance. *see* organizational governance
risk governance. *see* risk governance
risks, 176
Gramm-Leach-Bliley Act (GLBA), 260

H

Health Insurance Portability and Accountability Act (HIPAA), 85, 260

I

identity and access operations, 173
IEC 31010:2019, 59, **61**
implementation testing, 135–136
incident. *see also* risk events
management, 172, 216

incident response plan (IRP), 240
information security, 223–224
  availability, 225–226
  awareness training, 247–250
  business continuity management, 250–254
  confidentiality, 225–226
  controls, 236–237
  data privacy and data protection principles, 254–261
  frameworks, 226–233
  integrity, 225–226
  policies, 234–236
  practices, 237–247
  standards, 236
information technology (IT), 163–164
  data lifecycle management, 194–197
  disaster recovery management, 184–194
  emerging technologies, 207–215
  enterprise architecture (EA), 164–169
  operations management, 170–179
  project, program, and portfolio management, 179–184
  systems development life cycle (SDLC), 198–207
Infrastructure-as-a-Service (IaaS), 108, 179, 210
Infrastructure-as-Code (IaC), 184
infrastructure operations, 172
infrastructure risk
  cloud, 179
  databases, 178
  gateways, 177
  networks, 176–177
  operating systems, 177
  platforms, 172, 178–179
  software, 177–178
inherent risk, 83–85
insider threat. *see* people-related risks
integrated risk management (IRM), 8
integrity, 225–226
internal audit and accountability, 18
intrusion prevention systems (IPSs), 238
ISACA Code of Professional Ethics, xxviii
ISO 31000 ("Risk Management—Guidelines"), 16
ISO/IEC 20000, 171
ISO/IEC 27001/27002/27701/31000, 59, **61**, 128–129, 229–230
issues, findings, and exception management, 111
  action plan development, 115
  change management, 112
  closure and documentation, 116
  configuration management, 112
  determination of required action, 114
  identification and logging, 113–114
  implementation and tracking, 115
  initial triage and classification, 114
  reporting and continuous monitoring, 116
  review and approval, 114–115
  risk assessment and impact analysis, 114
  validation and verification, 115
IT Infrastructure Library (ITIL), 171
IT operations management
  backup, recovery, and data protection, 173
  change management, 172
  configuration management, 173
  core operations functions, 171–174
  device management, 172
  end-user services, 172
  identity and access operations, 173
  incident management, 172
  infrastructure operations, 172
  infrastructure risk, 176–179
  IT service desk, 172
  lifecycle and asset management, 174
  monitoring, logging, and observability, 173
  operational risk, 175–176
  operations governance and reporting, 174
  portfolio management, 179–184
  problem management, 172
  security management, 174
  vendor and third-party risks, 176

IT service continuity management (ITSCM), 184, *184*, **233–234**
IT service desk, 172

K

key control indicators (KCIs), 151
key performance indicators (KPIs), 147–149
key risk indicators (KRIs), 149–151

L

lagging KRIs, 150–151
leading KRIs, 150–151

M

machine learning (ML), 119–120
maturity assessment, 51
maximum tolerable downtime (MTD), 251–252
mitigation analysis, 69

N

National Institute of Standards and Technology (NIST), 126–127
network and systems security, 238–239
new risks, 32, 293
NIST Special Publication 800-30, *55–56*, **61**
non-repudiation, 243–244

O

The Open Group Architecture Framework (TOGAF), 166–167, *167*
Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE) methodology, *57–58*, **61**
operational management, 17
operational risk, *5*, 175
  environmental and physical risks, 176
  governance and compliance risks, 176
  people-related risks, 175
  process-related risks, 175
  security-related, 175
  technology-related risks, 175
  vendor and service provider coordination, 174
operations governance and reporting, 174
organizational culture, 6
organizational governance, 2–3
  business processes, 8–9
  organizational assets, 9–12
  organizational culture, 6
  policies and standards, 7–8
  stakeholders for cyber risk governance committee, 6
  strategy, goals, and objectives, 4
  structure, roles, and responsibilities, 4–5

P

Payment Card Industry Data Security Standard (PCI DSS), 127–128, **127**, 260
penetration test, 51
people-related risks, 175
physical access security, 245–246
Platform-as-a-Service (PaaS), 108–109, 179, 210
portfolio management, 179–184
prioritized recovery sequencing, 192
privileged access management (PAM), 245
problem management, 172
process-related risks, 175
professional ethics of risk management, 20–21
program management, 179–184
project drivers and risks, 180
project management, 179–184
project risk monitoring, 182

project threats, 180–181
project vulnerabilities, 181
protected health information (PHI), 85

Q

qualitative risks, 74, 280–282, 287
quantitative and qualitative analysis techniques, 71–75
quantitative risks, 71–73, 282, 287
quantum computing, 212

R

recovery point objective (RPO), 187–188, 251–252
recovery time objective (RTO), 187–188, 251–252
reference architectures, 165, 168
residual risk, 83–85
resilience engineering, 252–254
responsibility in third-party relationships, 104
risk acceptance, 101, 290–291
risk alignment, 165
risk analysis and evaluation
  artificial intelligence for, 86–87
  business impact analysis, 78–83
  inherent and residual risk, 83–86
  methodologies, 69–78
  precedence of prior risk treatment decisions, 288
  qualitative and quantitative risk, 287
  risk appetite and risk tolerance, 288
  risk ownership, 62–63
  risk ranking, 62
  risk register, 63, 64–66, 67, 68
  risk treatment options, selecting, 287–288
  worksheet, 286–287
risk and compliance management, 17
risk appetite, 18–19, 69, 288
risk assessment, 29–30, 188
  collecting evidence for, 51–54
  IT risk analysis and evaluation. *see* risk analysis and evaluation
  IT risk identification. *see* risk identification
  scope of, 51
  standards and frameworks, 55–61
  types of, 50–51
risk avoidance, 101–102, 292–293
risk capacity, 18
risk culture, 6, 19
risk discovery, 277–278
risk events, 34–37, *36*
risk governance, 13–14
  enterprise risk management, 15–16
  legal, regulatory, and contractual requirements, 20
  principal components of, **14**
  professional ethics of risk management, 20–21
  risk appetite, 18–19
  risk management frameworks, 15–16
  risk profile, 18
  risk tolerance, 18–19
  three lines of defense model, 17–18
risk identification, 31–33
  artificial intelligence for, 86
  brainstorming and, 34
  control deficiencies and gaps, 43–44
  control deficiency analysis, 40
  control effectiveness, existing, 41–42, 43
  control recommendations, 44–45
  risk and control ownership, 45
  risk events, 34–37
  risk scenario development, 47–49
  root cause analysis, 45–47
  threat landscape, 38–39
  threat modeling, 37, 38
  vulnerability analysis, 40
Risk IT Framework (ISACA), 15, 232–233, **233**
Risk Management Framework (RMF, NIST), 15–16, 226–228

risk management program, 271
  defined, 275
  foundations, 277
  lifecycle, 277–294
  purpose of, 276
  risk landscape, 272–275
risk mitigation, 100, 290, 294
risk monitoring and reporting
  actionable and decision-ready reports, 146
  audit and assessment results incorporation, 146
  communicating risk scenarios and business impact, 146–147
  data collection, aggregation, analysis, and validation, 142–143
  key control indicators, 151
  key performance indicators, 147–149
  key risk indicators, 149–151
  risk treatment plans, 139–141
  stakeholder-specific reporting needs, 145
  standard elements of risk reports, 145–146
  techniques, 144–145
risk ownership, 62–63
risk prioritization, 294
risk profile, 18
risk ranking, 62
risk registers, 63, **64–66**, 278–279, 293
  layers of, 282–283
  reviewing, 283–285
  sample entries in, **68**
  sources of information for, 67
  updating, 67
risk renewals, 294
risk response
  issues, findings, and exception management, 111–116
  management of emerging risk, 116–168
  risk and control ownership, 102–103
  risk treatment/risk response options, 99–102
  third-party risk management, 103–111
risk reviews, 293
risk scenario development, 47, 49
risk sharing, 100, 104
risk tolerance, 18–19, 69, 288
risk transference, 291–292
risk treatment, 294
  decision, 288, 290–293
  options, 99–102, 287–288
  plans, 139–141
  recommendation, 289
root cause analysis (RCA), 45–47

S

scenario identification, 188
schedule risk, 180
scope risk, 180
sectoral laws, 260–261
security architecture, 247
security incident response, 239–241
security information and event management systems (SIEM), 239
security management, 174
security orchestration and response (SOAR), 173
security-related operational risks, 175
semi-quantitative risk analysis, 74
Sherwood Applied Business Security Architecture (SABSA) model, 247
Software-as-a-Service (SaaS), 109, 179, 210
software composition analysis (SCA), 206
static application security testing (SAST), 206
stewardship, 195
strategic risk, 5
systems development life cycle (SDLC)
  code quality and secure coding standards, 205
  code review and peer validation, 206
  container and artifact security, 206
  continuous improvement and governance, 207
  design, 199–200

systems development life cycle (SDLC) (*Cont.*)
dynamic application security testing (DAST), 206
implementation, 200
logging, monitoring, and telemetry requirements, 207
operations and maintenance, 201
phases, 198–202
planning, 198–199
post-implementation review, 201
retirement, 201–202
risk factors, 203–204
roots of, 202
secrets and credential management, 207
secure development practices, 205–207
security requirements and design controls, 205
shifting left, 202–203
software composition analysis (SCA), 206
static application security testing (SAST), 206
testing, 200
threats, 204, 205
vulnerabilities, 204

T

tactical risks, *5*
tailgating, 246
technology-related risks, 175
third-party risk management (TPRM), 103, 176, 241
accountability/responsibility, 104
business process outsourcing (BPO), 111
cloud shared responsibility models, 107–110, *107*
continuous monitoring, 105–106
external ratings/continuous monitoring, 105–106
risk factors, 104–105
risk sharing, 104
threats, 106
vulnerabilities, 106–107
threat agents, 38, **39**
threat assessment, 37
threat intelligence platform (TIP), 173, 238
threat landscape, 38–39, **39**
threat modeling, 37, *38*, 50
three lines of defense model
internal audit and accountability, 18
operational management, 17
risk and compliance, 17
top-down risk scenario development, 48, *48*
training programs
audience-specific content, 248
behavioral reinforcement, 249
delivery and measurement, 249–250
leadership-level risk training, 249

U

U.S. State Breach Notification Laws, 261

V

vendor and service provider coordination, 174
vendor and third-party risks, 176
vulnerability
analysis, 40, *69*
assessment, 40, 50
management, 239
scanner, 40

Z

Zachman Enterprise Architecture Framework, 168, **169**
zero-trust architecture (ZTA), 34, 214

Online Test Bank

To help you study for your CRISC Certified in Risk and Information Systems Control Exam, register to gain one year of FREE access after activation to the online interactive test bank—included 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:

- Practice test questions, so you can practice in a timed and graded setting
- Flashcards
- A searchable glossary

SYBEX®
A Wiley Brand

Keep Learning and Earning!

...with certification study guides by Peter H. Gregory

CISA Study Guide

https://www.amazon.com/dp/1394288387/

AIGP Study Guide

https://www.amazon.com/dp/139436394X/

CISM Exam Guide

https://www.amazon.com/dp/1264268319/

CISSP For Dummies

https://www.amazon.com/dp/1394261772/

Amazon Author Page

https://www.amazon.com/stores/Peter-H.-Gregory/author/B000APLUN6

PETER H. GREGORY

CISSP, CISM, CISA, CRISC, CDPSE, CIPM, CCSK

Peter is a career cybersecurity leader and practitioner with 30 years of real-world security engineering, risk, privacy, and leadership experience. He has written over 50 books on cybersecurity and privacy certifications, as well as training courses on O'Reilly. He also has 10 years of experience teaching cybersecurity at the University of Washington. www.peterhgregory.com



WILEY END USER LICENSE AGREEMENT

Go to www.wiley.com/go/eula to access Wiley’s ebook EULA.