Certificate of Cloud Auditing Knowledge™

Study Guide

CCAK™

Certificate of Cloud Auditing Knowledge

A Cloud Security Alliance® and ISACA® Credential

CSA cloud security alliance®

ISACA®

About Cloud Security Alliance

The Cloud Security Alliance® (CSA) (www.cloudsecurityalliance.org) is the world's leading organization dedicated to defining and raising awareness of best practices to help ensure a secure cloud computing environment. Cloud Security Alliance harnesses the subject matter expertise of industry practitioners, associations, governments, and its corporate and individual members to offer cloud security-specific research, education, certification, events and products. Cloud Security Alliance activities, knowledge and extensive network benefit the entire community impacted by cloud—from providers and customers, to governments, entrepreneurs and the assurance industry—and provide a forum through which diverse parties can work together to create and maintain a trusted cloud ecosystem. Twitter @cloudsa.

About ISACA

For more than 50 years, ISACA® (www.isaca.org) has advanced the best talent, expertise and learning in technology. ISACA equips individuals with knowledge, credentials, education and community to progress their careers and transform their organizations, and enables enterprises to train and build quality teams that effectively drive IT audit, risk management and security priorities forward. ISACA advances and validates business-critical skills and knowledge through the globally respected Certified Information Systems Auditor® (CISA®), Certified in Risk and Information Systems Control® (CRISC®), Certified Information Security Manager® (CISM®), Certified in the Governance of Enterprise IT® (CGEIT®) and Certified Data Privacy Solutions Engineer™ (CDPSE™) credentials. ISACA is a global professional association and learning organization that leverages the expertise of more than 150,000 members who work in information security, governance, assurance, risk and privacy to drive innovation through technology. It has a presence in 188 countries, including more than 220 chapters worldwide. In 2020, ISACA launched One In Tech, a philanthropic foundation that supports IT education and career pathways for under-resourced, under-represented populations.

Disclaimer

Cloud Security Alliance and ISACA have designed and created *Certificate of Cloud Auditing Knowledge™ Study Guide* (the "Work") primarily as an educational resource for auditing, security and governance professionals. Cloud Security Alliance and ISACA make no claim that use of any of the Work will assure a successful outcome. The Work should not be considered inclusive of all proper information, procedures and tests or exclusive of other information, procedures and tests that are reasonably directed to obtaining the same results. In determining the propriety of any specific information, procedure or test, professionals should apply their own professional judgment to the specific circumstances presented by the particular systems or information technology environment.

Reservation of Rights

Cloud Security Alliance
2212 Queen Anne Ave N
Seattle, WA 98109, USA
Phone: +1.360.746.2689
Fax: +1.206.832.3513
Contact us: support@cloudsecurityalliance.org
Website: https://cloudsecurityalliance.org/

Provide Feedback: support@cloudsecurityalliance.org
CSA Circle Online Community:
https://circle.cloudsecurityalliance.org/

Twitter: https://twitter.com/cloudsa
LinkedIn: www.linkedin.com/company/cloud-security-alliance
Facebook: www.facebook.com/csacloudfiles

ISACA
1700 E. Golf Road, Suite 400
Schaumburg, IL 60173, USA
Phone: +1.847.660.5505
Fax: +1.847.253.1755
Contact us: https://support.isaca.org
Website: www.isaca.org

Provide Feedback: https://support.isaca.org
ISACA Online Forums:
https://engage.isaca.org/onlineforums

Twitter: http://twitter.com/ISACANews
LinkedIn: www.linkedin.com/company/isaca
Facebook: www.facebook.com/ISACAGlobal
Instagram: www.instagram.com/isacanews/

ISBN 978-1-60420-863-4
Certificate of Cloud Auditing Knowledge™ Study Guide
Printed in the United States of America

Acknowledgments

Cloud Security Alliance Cloud Audit Expert Group

The *Certificate of Cloud Auditing Knowledge™ Study Guide* was developed with the support of the Cloud Security Alliance Cloud Audit Expert Group, whose lead developers, editors and expert reviewers include contributors and volunteers from across the globe, reflecting a wide variety of industries—both cloud service consumers and providers. The Cloud Audit Expert Group surveyed cloud ecosystem stakeholders—from cloud service providers and cloud customers to auditors and regulators—and assessed relevant job functions (e.g., board of directors, CEO, CTO, CISO, CIO, privacy office, legal, internal audit, procurement, IT, security and development teams) and, from cumulative stakeholder input, established the value proposition, scope, learning objectives and curriculum of the *Certificate of Cloud Auditing Knowledge Study Guide*.

Lead Developers

Craig Balding, CISSP, CITP, MBCS, Resilient Security Ltd, UK (Chapters 1, 5)
Doug Barbin, CCSK, CISSP, CPA, QSA, ISO 27001 Lead Auditor, Schellman & Company LLC, USA (Chapters 2, 6)
Jon-Michael Brook, CCSK, CISSP, CSA Research Fellow, Starbucks, USA (Chapter 2)
Vince Campitelli, Cloud Security Alliance, USA (Chapter 6)
Daniele Catteddu, CISM, Cloud Security Alliance, Italy (Chapters 2, 3, 5, 6)
Moshe Ferber, CCSK, CCSP, Israel (Chapter 1)
David Frei, CIPP, CISSP, USA (Chapter 2)
Francoise Gilbert, CIPM, CIPP/E, CIPP/US, DataMinding, USA (Chapter 1)
John Guckian, CCSK, CISSP, ITIL, USA (Chapters 2, 5)
Jim de Haas, ABN AMRO Bank, Netherlands (Chapter 2)
Matthew Hoerig, AWS CLF-C01, AWS SAA-C01, CCDP, CCIP, CCNP, CCSK, CCSP, CISSP, VCE-CIIEc, VCP 6, VCP-NV, Canada (Chapter 2)
Murat Lostar, Bugbounter, Turkey (Chapter 1)
Harry Lu, CCSP, PwC, USA (Chapter 7)
John Di Maria, AMBCI, CERP, CSSBB, HISP, MHISP, Cloud Security Alliance, USA (Chapters 6, 9)
Rich Mogull, DisruptOps and Securosis, USA (Chapter 8)
Jacques Nack Ngue, CIPP, CISSP, USA (Chapter 2)
Alain Pannetrat, Ph.D., Cloud Security Alliance, Greece (Chapter 2)
Michael Roza, CISA, CIA, CPA, EMBA, EVC, USA and Belgium (Chapter 4)
Damir Savanovic, CISA, CISM, Spain (Chapter 2)
Andrew Williams, Coalfire, USA (Chapter 6)
Steven Woodward (Chapter 2)

Contributing Editors

Ryan Bergsma, CCSK, Cloud Security Alliance, USA
Suhas Bhat
Felipe Castro, CRISC, Hewlett Packard Enterprise, Mexico
Willibert Fabritius, BSI Group, USA
Kevin Fumai, USA
Ricci SC Leong, Ph.D., CISA, CISM, CCSK, CCSP, CEH, CISSP, GPEN, eWalker Consulting (HK) Ltd, Hong Kong
Peter J. Knuz, SGL Carbon, Germany
Robin Lyons, CISA, CDPSE, CIA, USA
Ganesh Mahalingam, CCSK, CISA, AWS SAA, CySA+, ITIL, PMP, SAFe POPM, Towson University, USA
Steven Mezzio, Ph.D., CISA, CISSP, CPA, FSA I, Pace University, USA
Josephine Jackeline Naicker, CISA, CCSP, CFE, CIA, CISSP, CPA, CRMA, ATCO, Canada
Adalberto Afonso Navarro Valle, CISA, CDPSE, Security+, CobiT, ITIL, RSM US, USA
Gary Nelson, CISA, CISSP, CPA, FIP, Schellman & Company, USA

Acknowledgments *(cont.)*

Gordon Smith, Ireland
Venky Thiruvengadam, CISM, Cybersecurity Audit, CSX-P, SSCP, Canada
Carlos E. Victoria, CISA, CCSK, CISSP, USA
Steven Woodward, CCSK, CFPS, CSQA, Cloud Perspectives, Canada

Expert Reviewers

Sanjiv Agarwala, CISA, CISM, CGEIT, CDPSE, CISSP, FBCI, Oxygen Consulting Services, India
Brandy Alexander, CISA, CISM, CCSP, QSA, CyberGuard Compliance, USA
Anita Ali, CISA, CPA, Law Society of Ontario, Canada
Renee Bernadette Bergeron, CISA, CISM, Thomson Reuters, USA
Anil Chiplunkar, CISM, CAME, CDPP, CFAP, CFE, Principal Auditor ISO 27001, India
Avani Desai, CISA, CCSK, CIPP, CISSP, CPA, Schellman & Company, LLC, USA
Givemore Dube, CISA, CRISC, CISM, CGEIT, CDPSE, CCSP, CISSP, EY, Zimbabwe
Peter HJ van Eijik, CCSK, CCSP, ClubCloudComputing and Hogeschool Utrecht, Netherlands
Marco Ermini, Ph.D., CISA, CRISC, CISM, CDPSE, AWS-CCP, AWS-CSAA, CCS-DCD, CIPM, CIPP/E, CISSP, GCIH, ITILv3, Orange Business Services, Germany
Eyal Estrin, CISA, CISM, CCSK, CCSP, CISSP, Israel
David A. Gittens, CISA, CRISC, CISM, CDPSE, CSX-P, AMBCI, CCSK, CCSP, CHFI, CISSP, eCPPT, eNDP, HISP, Security+, Barbados
Eric J. Gonzalez-Mendez, C)DFE, C)PEH, C)PTE, ISO 27001 Lead Auditor, LSSGB, PMP, Hewlett Packard Enterprise, Puerto Rico
Alison Gunnels
Arif Hameed, CISA, CRISC, ABCP, CISSP, ISO 27001 LA, Equifax Canada, Canada
Rupert Hargrave, CISA, CCSK, Deloitte LLP, United Kingdom
Roberto Hernandez Rojas Valderrama, CISA, CRISC, CISM, CGEIT, CDPSE, COBIT 2019-F, CSX-F, ISO 27001 LA, ITIL-F, PMP, Mexico
Eric Hibbard, CISA, CCSK, CIPT, CISSP-ISSAP, ISSEP, ISSMP, USA
Dietmar Hinkel, CISA, Germany
Anand Jha, CISA, CDPSE, Azure Solution Architect, CCSK, CISSP, ISO 27001 LA, GCIH, GPEN, India
Leighton Johnson, CISA, CCSK, CISSP-ISSEP, ISFMT, ISO 27001 LA, USA
Shamik Kacker, CRISC, CISM, CCSK, CCSP, CISSP, ITIL Expert, TOGAF 9 Foundation, Dell Technologies, USA
Ramaswami Karunanithi, CISA, CRISC, CGEIT, CSX-A, CA, CAMS, CBCI, CCSK, CCSP, CFE, CFSA, CGAP, CGMA, CHFI, CIA, CIPP/E, CISSP, CMA, CPA, CPRM, CRMA, CSCA, CSFX, FCS, GRCA, GRCP, Lead Auditor ISO 27001, PBA, PMP, Prince2 Practitioner, RMP, New South Wales Government–Legal Aid, Australia
Fahima Khan
Mohammed J. Khan, CISA, CRISC, CDPSE, CIPM, Baxter, USA
Bilal Khattak, CISA, CRISC, CISM, CGEIT, Volkswagen Financial Services AG, Germany
Monica LaCroix, CISA, CRISC, CISSP, ISO Lead, MCSE, QSA, Coalfire, USA
A.T. Manjunath, CISA, CSX-A, CSA STAR Auditor, CSA-CCSK, CISSP, SSGB, Applied Materials, India
Luciano Moreira da Cruz, CCSK, CSA STAR Auditor, DSOE, MCSE, MVP, AL ISO 27001, 9001, 27017, 27018, 2000, Cloud Legion, Argentina
Matt Nixon, CISA, CRISC, CISM, CISSP, GSEC, Singapore
Emmanuel Nyeanchi, Ph.D., CISM, USA
Mark Pearce, CISA, CRISC, CISM, CGEIT, CDPSE, CIIP, CIPM, GSLC, IRMCert, Secureworks, UK
Sridhar Pippari, CISA, CISM, CCSK (v3 & v4), PSM (I), SANS GCIH, India
Kurt Siefried, CCSK, Cloud Security Alliance, Canada
Eleftherios Skoutaris, Cloud Security Alliance, Greece
Matt W. Snider, CISA, CSX-P, CISSP, Denim Group, USA

Acknowledgments *(cont.)*

Fadi Sodah, CISA, CISM, CISSP, Amman Stock Exchange, Jordan
Annie Sokol, NIST, USA
Srinivasan Subramaniam Muthukondapuram, CISA, CISM, CDPSE, CEH, CISSP, CREST-CPSA, OSCP, ECSA, Auriseg Consulting Private Limited, India
Han Sang Sung, CCSK, KPMG (Korea), South Korea
Stacy Thomas, CISA, CRISC, CIA, CIPP-US, CISSP, USA
Manoj Wadhwa, CISA, CRISC, CISM, EY, Singapore
Dominique West, CISSP, Security in Color, USA
Xiaoqi Zhao, CISA, CISSP, ITIL, PMP, Volvo Group, China

Cloud Security Alliance Board of Directors

ISACA Board of Directors

Page intentionally left blank

Table of Contents

Chapter 2:
Cloud Compliance Program....105

Chapter 3:
Introducing CCM and CAIQ 179

Chapter 4:

Chapter 5:

Cloud Auditing 251

Chapter 6:

LIST OF FIGURES

Chapter 1. Cloud Governance

Chapter 2. Cloud Compliance Program

Chapter 3. Introducing CCM and CAIQ

Chapter 4. A Threat Analysis Methodology for Cloud Using CCM

Chapter 5. Cloud Auditing

Chapter 6. Evaluating a Cloud Compliance Program

Chapter 7. CCM Auditing Guidelines

Chapter 8. Continuous Assurance and Compliance

Chapter 9. Security Trust Assurance and Risk (STAR) Program

Introduction

The *Certificate of Cloud Auditing Knowledge™ Study Guide* is designed to help individuals acquire foundational knowledge of cloud governance, security and auditing, and supports students in preparing to qualify for the Cloud Security Alliance® (CSA) and ISACA® Certificate of Cloud Auditing Knowledge™ (CCAK).

Topics in this guide are presented at a foundational level and include basic terminology, concepts, examples, general practices and explanations of cloud security principles. The *Certificate of Cloud Auditing Knowledge Study Guide* provides the IT or business professional with knowledge in the following main cloud content areas:

- Types of security
- Security governance
- Business continuity and disaster recovery
- Recovery
- Privacy vs. security
- Threat landscape
- Cyberrisk
- Cyberattacks
- Securing assets
- Security operations and response
- Vulnerability management
- Penetration testing
- Incident response
- Monitoring

Every chapter contains the following:

- **Learning objectives**—Knowledge that learners should understand after reading the chapter
- **Sections**—Knowledge content
- **Knowledge checks**—Quiz on the chapter knowledge

Note: This guide primarily references two entities: the cloud service provider (CSP) and its customer, which is sometimes called the cloud service consumer, (CSC). The terms organization, customer and CSC are used interchangeably in this guide to refer to the entity that buys and uses a cloud service from a CSP. The term client refers to the end customer or user of the organization product or service. In chapter 7, which covers audits of CSPs and CSCs, the term organization may refer to the CSP customer and CSP—the distinction is clear from the context.

Page intentionally left blank

Chapter 1:

Cloud Governance

Cloud Governance

Cloud Governance

1.1 Learning Objectives

After completing this chapter, learners will be able to:

1. Describe cloud governance concepts.
2. Explain cloud trust, transparency and assurance.
3. Identify cloud governance frameworks and requirements.
4. Discuss cloud risk management and cloud compliance considerations.
5. Distinguish cloud governance tools and uses.

1.2 Overview

This chapter introduces the concept of cloud governance, explains why it is important for an effective security and compliance posture, and describes the relationship between governance and IT security. Cloud use and deployment models are examined, while considering their impact on traditional approaches to governance. This chapter also covers the place of cloud governance in overall corporate governance and introduces the shared responsibility model.

Further, this chapter covers cloud trust, transparency and assurance; a cloud governance framework; and governance requirements that are specific to cloud computing. Examples of cloud risk and approaches to identifying and managing it are included.

The sections about cloud governance tools address the roles of policy enforcement, contracts and security assessments, and provide an overview of audit types that are applicable to the cloud environment.

1.3 Cloud Governance Overview

This section introduces and provides an overview of cloud governance.

1.3.1 What Is Governance?

ISACA defines governance as the process by which an organization ensures that stakeholder needs, conditions and options are evaluated to achieve balanced, agreed-upon objectives. Governance is based on a framework of policies, procedures and controls designed to promote transparency and accountability to defined standards.

Strong governance practices address strategic guidance, risk management and mitigation, compliance monitoring and remediation, budget allocation, and cost controls. IT governance ensures that information-related technologies support and enable organization strategies and objectives. Organizations with mature IT governance practices can measure their results against their business objectives and make ongoing enhancements.

For more information on governance standards, see:

- ISO/IEC 38500:2015 *Information Technology – Governance of IT for the organization*
- COBIT® 2019
- ISO/IEC 27014:2013 *Information Technology – Security techniques – Governance of information security*
- The Open Group® Cloud Computing Governance Framework

Examples of laws and regulations that affect IT governance requirements include:

- Gramm-Leach-Bliley Act (GLBA)
- Sarbanes-Oxley Act
- The Health Insurance Portability and Accountability Act (HIPAA)
- General Data Protection Regulation (GDPR)

1.3.2 Importance of Governance to Security

Governance plays a vital role in establishing accountability and providing oversight to ensure the overall security posture of an organization. Governance provides a defined approach to IT security that can be applied consistently across the many IT service delivery models. Governance aligns security goals with business objectives through a risk-based security approach.

Most organizations have a complex mix of IT services that are provided in different ways—on premises and managed internally or by a third party, and traditional hosting and cloud services delivered by a third party. Governance provides a way for organizations to ensure that they can direct, measure, compare and enforce IT security and compliance consistently across all delivery models.

Relationship Between Governance and IT Security

Governance can help organizations identify and measure all types of risk, such as financial, credit, legal, regulatory and reputational risk. These measurements enable the creation and evolution of new and better ways to manage and mitigate risk.

Governance has an especially deep relationship to IT and its risk categories, which are usually described as confidentiality, integrity and availability risk to information, and are grouped under the information risk management title (**figure 1.1**).

Figure 1.1—Hierarchy of Governance and Information Risk

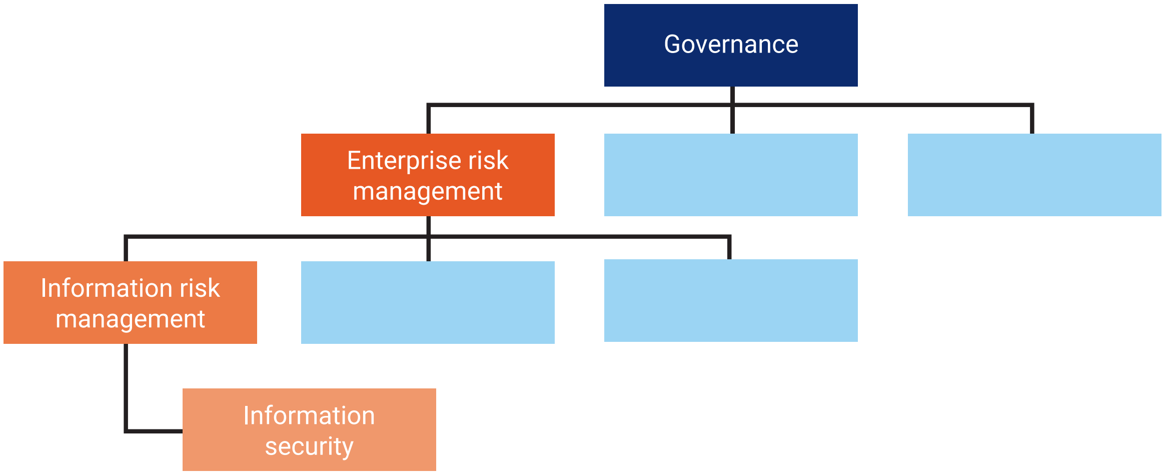

Source: Cloud Security Alliance CCSK training

The Importance of Governance for IT Security

Governance practices are critical for aligning business goals with IT resources and road maps, i.e., part of governance is establishing the organization perception of risk and defining its risk appetite. This allows an organization to establish governance policies and business processes that enable it to operate in a deliberate, disciplined manner, and to identify policy breaches and related risk—fundamental to any IT security policy. Without proper governance, it is difficult for the security professional to handle risk and implement controls in a way that supports business objectives.

Example: Startups often choose to work with cloud providers that are less mature but provide good value for money, whereas organizations with established governance policies look to more mature CSPs that make significant investments in security as a core component of the cloud service and, thereby, offer stronger protection for data and other critical assets.

1.3.3 Cloud Use Models and Potential Impact on Traditional Governance Models

Cloud computing heavily affects governance, because cloud introduces a new business model, new technologies that require unfamiliar types of controls and processes, and new third parties into the IT ecosystem. Any of those factors might, by themselves, introduce governance changes. The combination of all three requires organizations and their IT functions to realign their model to reflect the new reality.

Complexities of Cloud Governance

Cloud computing can introduce many complexities into the organization governance approach. Although some considerations may vary depending on the specific type of cloud service being deployed (e.g., SaaS, IaaS, or PaaS), the following considerations apply generally to most cloud services:

- Cloud might impose a loss of direct control over the IT infrastructure, forcing an organization to adopt a new governance framework and process.
- Cloud services and data may span multiple jurisdictions, forcing customers to comply with more laws and regulations.
- Visibility and transparency into some cloud services can be challenging.
- Using cloud solutions does not mean outsourcing the organization accountability of its controls to a third or fourth party.
- Data ownership rights might not be intuitively clear and need careful examination.
- Most providers have a standard offering that cannot be customized according to a customer’s specific requirements.
- Cloud providers might demonstrate different levels of maturity and variety of services, licenses and models, which complicates adoption of a one-size-fits-all cloud policy.
- Cloud services are often built on a chain of providers, which makes scoping of governance activities challenging (e.g., a SaaS provider that is running on the infrastructure of an IaaS provider).
- Cloud uses a shared responsibility model, which requires specific mapping of controls and responsibilities between provider and customer.
- Hybrid cloud models can complicate governance due to complexities of producing clear boundaries between provider responsibilities and customer responsibilities.
- In cloud services, customers must rely more on assessment than actual testing.

- Cloud providers change rapidly, which puts long-term planning and complex policies at risk of no longer being relevant within a very short time.

Shared Responsibility Model and Reality (No.1)

The shared responsibility model establishes that an organization and CSP have separate but complementary obligations to ensure that the cloud service operates and remains protected as intended. In reality, this sharing of responsibilities often extends beyond the CSP to agents it uses to deliver the service and other parties—such as a cloud platform integrator, or software development companies that create or manage applications running on top of the cloud platform.

Impact of Cloud Service Models

Following are the cloud service models and descriptions of their impact on the traditional governance models:

- **Infrastructure as a Service (IaaS)**—Similar to traditional data center security and traditional hosting and outsourcing governance frameworks, IaaS can work with minimal adjustments. In IaaS, the customer has a significant share of the responsibilities, and the service tends to be more static than other cloud services, making cloud governance for IaaS an easier task. However, IaaS providers are usually large-scale providers that can be reluctant to customize their offerings, assume new obligations or significantly change their contractual terms for a specific customer.

  In building cloud governance programs for IaaS, most of the responsibility lies with the customer, and the sharing of responsibilities between provider and customer is reasonably defined for mature providers. IaaS governance includes achieving technical competency and building usage policies that allow freedom of innovation, yet are balanced against defined organizational limitations.

  With IaaS, organizations typically limit the number of providers to simplify their operations and better manage their environments. Therefore, it is highly beneficial to create and enforce a central policy across all vendors and all key security domains (e.g., business continuity and backup, identity, privileges and access management, logging and monitoring, and compliance).
- **Platform as a Service (PaaS)**—PaaS offers cloud customers the environments and tools to develop, run and manage their own applications. PaaS is an evolving offering with different capabilities depending on the provider. However, it usually presents thin boundaries between vendor and customer responsibilities, which complicates cloud governance. The governance model for PaaS requires careful consideration of the contractual terms established with each provider to create a harmonized, holistic governance policy to enforce.

  Since IaaS and PaaS are usually combined, cloud governance recommendations about IaaS are also relevant to PaaS. Further, the adoption of PaaS usually involves a transition to software development and IT operations (i.e., DevOps) practices, which usually leads to different approaches to cloud governance policies. Whether an organization embraces DevOps or not has substantial repercussions for its cloud governance model.

  When it comes to cloud governance, the most important factors to consider are who is setting the requirements and who is driving the decisions. Is the DevOps team or the infrastructure team taking the lead? Who are the stakeholders involved in the decision-making process? Are organization-wide standards applied, or are there cloud-specific exceptions? What are the trade-offs?

  Other PaaS governance recommendations that must be considered follow:

  - Metrics to measure security
  - Achievement of technical competency
  - Policies and controls for developers and DevOps

PaaS governance policies should function like guardrails—allow enough freedom to innovate, but make sure that certain security requirements are always met.

- **Software as a Service (SaaS)**—SaaS presents different challenges from IaaS and PaaS, which organizations usually implement in small numbers and are centralized by nature. Typical large-scale organizations can use hundreds of decentralized SaaS applications. It is not uncommon for global organizations to deploy different SaaS providers in different geographical locations.

  SaaS platforms usually present a wide variety of services to the market with different maturity levels. Providers can range from software giants with unified and mature services to niche providers that may be willing to customize their services for each customer. Therefore, developing a centralized, one-size-fits-all SaaS governance model is challenging.

  Cloud customers need to consider the natural evolution of the adoption of certain SaaS services. An organization may begin a specific SaaS service journey with limited use. However, the actual scope of usage can grow over time and produce new risk that was not evaluated during the early stages.

  SaaS is the cloud service model with the least visibility and control into the underlying infrastructure. Customers need to rely more on assessment, contracts, SLAs (service level agreements) and monitoring to provision controls and ensure that the services being provided are working effectively.

  SaaS governance models should require involvement of the business departments that can best evaluate the real impact of risk for a specific SaaS service. Organizations must apply existing security policies and standards to SaaS applications and adapt to the unique characteristics of each cloud. Because of the variety of SaaS applications, there are also different risk levels. Governance processes, policies and monitoring should address the different risk levels.

  Other SaaS governance recommendations are:

  - Create a cloud version of the contract template and add a cloud-specific delivery agreement with the technical details on security
  - Establish a central cloud register that lists all SaaS applications in use

A cloud governance program should cover all aspects of the various cloud services models, because the boundaries between service models are very thin, and a cloud migration program can include workloads on any given cloud service model (IaaS/PaaS/SaaS).

Shared Responsibility Model and Reality (No.2)

Textbooks discuss how responsibility shifts from customer to provider with a move from IaaS to SaaS, but, in reality, the borders between IaaS, PaaS and SaaS are fading fast. With the click of a button, a developer can move from one service model to another, and an employee can purchase a new cloud service outside the traditional procurement function. When auditing large cloud providers, all services cannot be categorized as IaaS or PaaS; some services are relevant to both.

Impact of Cloud Deployment Models

Following are the cloud deployment models and descriptions of their impact on the traditional governance models:

- **Public cloud**—Public cloud is the most popular cloud deployment model. Public cloud providers deliver unified standard services for all customers and, therefore, might object to requests to customize the service. This complicates cloud governance, because providers are responsible for governing their own infrastructure, services and employees. Heavy reliance on customer configurations to public cloud services produces governance challenges initially and over time.

The value of public cloud value offering relies on multi-tenancy, which brings further governance challenges, such as segmentation and isolation. Certain actions, such as security scanning and penetration testing, are prohibited or limited, and visibility into the underlying infrastructure is limited or nonexistent.

These characteristics of the cloud present opportunities for customers to adopt new ways of governance and new models that rely on data provided through vendor risk management, service level agreements (SLAs), third-party audits, and compliance reports based on laws and regulations. The effectiveness of this new approach is measured by the cloud consumer's ability to mitigate the unique risk that cloud computing introduces.

- **Private cloud**—Private clouds can be owned, managed or hosted by the organization or by a third party. Cloud governance of self-managed private clouds is the closest to traditional IT governance, but it should also address the cloud platform attack vectors, multi-tenancy at the organizational level and unique cloud capabilities, such as automation.

  Governance of private clouds managed by third parties is closest to the traditional outsourcing models. Governance challenges include understanding the shared responsibility matrix, setting appropriate SLAs, and building third-party capabilities to monitor for policy violations, malicious insiders and rogue employees. One of the biggest challenges of private cloud governance is ensuring that the platform and services are kept up-to-date.

- **Hybrid cloud**—Hybrid cloud services typically span two cloud deployment models—private and public. Hybrid clouds can be implemented in several ways that have very few things in common, which complicates establishment of policy guidelines and the shared responsibility model.

  Hybrid cloud governance challenges include aligning the SLA and shared responsibility model between provider and customer, protecting the internal perimeter, scaling the security configuration, and addressing gaps in cloud skill sets and maturity.

- **Community cloud**—Community cloud encompasses a wide range of services that have very little in common. Usually, community cloud involves third-party management and hosting of the cloud service. Many organizations share the platform itself, but it is not fully public, which helps to mitigate the multi-tenancy challenges.

  Cloud governance challenges include identifying the relevant stakeholders and building the correct shared responsibility model. Relationships among organizations using the same community cloud and anticipated risk from multi-tenancy issues require specific governance attention.

1.3.4 How Cloud Governance Fits Into Overall Corporate Governance

Governance and IT Governance

Cloud governance is part of the overall corporate governance policy that organizations should implement. As part of an overall risk management strategy, organizations should define the principles of cloud governance, including decision criteria, minimal security requirements, and reporting and monitoring capabilities.

Organizational governance policy determines the organization tolerance for risk. IT governance policy identifies the right stakeholders and sets the control framework. Cloud governance relies on those determinations as a baseline for cloud-specific policy.

Figure 1.2 is an example of mapping between COBIT 2019 principles and cloud governance policy.

Figure 1.2—Mapping of COBIT 2019 Principles and Cloud Governance

COBIT 2019 Principles	Cloud Governance Considerations
Meeting stakeholders' needs	Part of cloud governance is identifying the stakeholders for cloud migration and determining who benefits and who is accountable for the risk.
Covering the enterprise end to end	Cloud governance scope includes cloud activities and identifying the relevant resources in the organization.
Applying a single integrated framework	Cloud governance maps a cloud governance standard, e.g., CSA STAR, Cloud Controls Matrix (CCM) and CSA Enterprise Architecture.
Enabling a holistic approach	Cloud governance identifies the enablers for cloud migrations, e.g., team members with relevant skill sets.
Separating governance from management	Cloud governance sets the objectives of cloud migration and the metrics for monitoring it.

Foundation of Cloud Governance

Cloud governance policy reflects the overall risk appetite of the organization. Cloud governance addresses the process of cloud migration—including such things as provider evaluation criteria and security requirements. **Figure 1.3** highlights certain decisions to consider when building cloud governance policy:

Figure 1.3—Cloud Governance Policy Considerations

Cloud Governance Policy	Considerations
Types of data processes that are allowed to migrate into the cloud	Laws and regulations and business critically determine the types of data processes permitted.
Locations and jurisdictions to maintain data	Privacy laws can complicate cross-border transfers.
Terms under which data can migrate	Certain applications can only migrate after formal board approval.
Relevant stakeholders	Cloud migration involves many stakeholders, e.g., legal, IT and procurement.
Minimal requirements for cloud providers	CSP must align with security industry best practices. e.g., CSA CCM, ISO27001 and 27017.
Mandatory controls to be implemented	All personal data and confidential information should be encrypted for pseudonymized or anonymized.
Monitoring and reporting	Requirements for customer visibility into the service must be included.

Tip: Certain regulations worldwide require board or senior management approval for certain cloud migrations, e.g., the Israeli central bank regulation requires that cloud policy must be approved by the board of directors of the financial institute.

Governance and the Shared Responsibility Matrix

Cloud governance relies on the shared responsibility model to identify the division of responsibilities between the provider and the customer. The result of such a process of responsibility attribution is a control matrix that identifies the associated person or team responsible for implementation. The shared responsibility control matrix is ideally reflected in the instructions of the contract signed with the cloud provider and should be recorded in a document that is agreed on by both the CSP and customer. **Figure 1.4** shows the path from governance to risk management:

Figure 1.4—Path From Governance to Risk Management

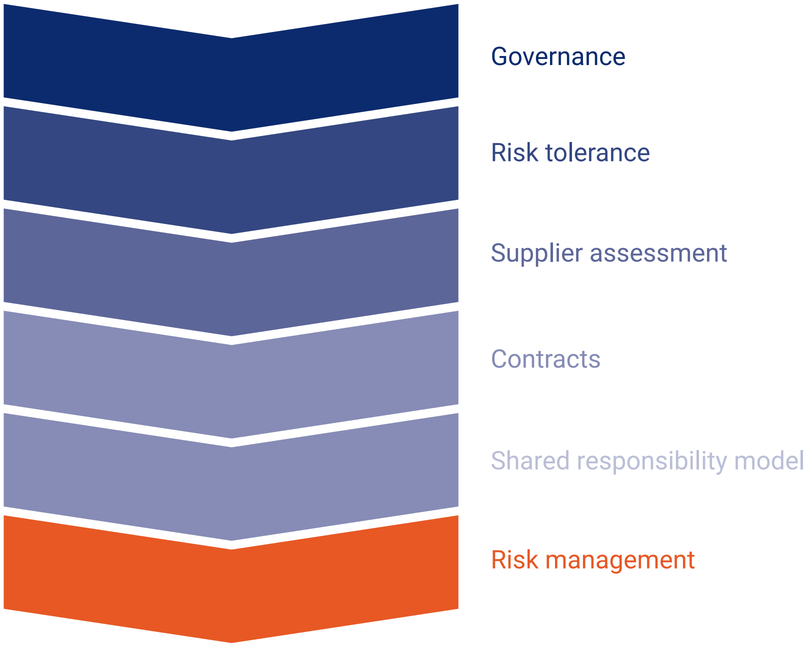

1.3.5 Types of Accountability/Organizational Models

Section 1.4.7 covers the importance of accountability to the governance process. This section focuses on the organizational roles, stakeholders and structure that are vital for ensuring proper accountability of cloud operations.

The IT governance framework (ISO/IEC 38500) provides the following guiding principles for good corporate governance of IT:

- Responsibility
- Strategy
- Acquisition
- Performance
- Conformance
- Human behavior or organization culture

When implementing cloud governance practices, it is important to keep these six guiding principles in mind and to establish the following roles:

- **GRC department**—Responsible for identifying risk and advising on mitigating controls
- **Legal department**—Responsible for contract management and specific issues, such as compliance with applicable laws and regulations (to unify and broaden the concepts of cross-border data flows and privacy requirements)
- **Procurement department**—Responsible for creating requests for proposals/information/quotations (RFPs/RFIs/RFQs), and vendor due diligence, selection, negotiation and oversight
- **IT/operations**—Responsible for building the right architecture

Because cloud adoption requires the attention and involvement of multiple stakeholders, organizations tend to manage cloud strategy and migration within a cloud committee or cloud center of excellence (CCOE).

To set the various responsibilities, a responsible, accountable, consulted and informed (RACI) matrix is often used to identify the various roles and responsibilities within the cloud adoption process. It is used for clarifying and defining roles and responsibilities in cross-functional or cross-departmental projects and processes.

The RACI roles:

- **Responsible**—Those who do the work to complete the task. There is at least one role with a participation type of responsible, although others can be delegated to assist in the work required.
- **Accountable**—The one ultimately answerable for the correct and thorough completion of the deliverable or task; the one who ensures the prerequisites of the task are met and delegates the work to those responsible. An accountable must sign off on (approve) the work that those responsible provide. There must be only one accountable specified for each task or deliverable.
- **Consulted (sometimes consultant or counsel)**—Those whose opinions are sought, typically subject matter experts, and with whom there is two-way communication.
- **Informed (also informee)**—Those who are kept up to date on progress, often only on completion of the task or deliverable, and with whom there is just one-way communication.

Figure 1.5 is an example of a RACI matrix describing cloud migration roles and responsibilities.

Figure 1.5—Example Cloud Migration RACI Matrix

	GRC Department	Legal Department	Security Team	IT Operations
Setting the provider requirements	R/A	C	C	I
Building governance scheme	R/A	C	C	I
Cloud vendor assessment	A	I	R	R
Building the architecture	I	I	A/R	R
Actual cloud migration	I	I	C	A/R

1.3.6 Cloud Governance Scope and the Gaps Created by New Hybrid Technology Models

In terms of scope, cloud governance reaches beyond typical security considerations. Following are some other aspects of cloud governance:

- **Cost management**—Every action on the cloud is charged, and lack of oversight and control in this area can result in increasing costs.
- **Compliance**—Requirements can arise from multiple local and international sources (based on the organization locations), such as laws, regulations, industry guidelines/requirements and business decisions.
- **Resilience**—Cloud applications can be resilient, with redundancy capabilities, but they can also prove not to be. Developing, testing and implementing a disaster recovery and business continuity plan that incorporates cloud services requires ongoing time and effort.
- **Portability and interoperability**—Portability is the ability to move data and workloads from one cloud provider to another to avoid vendor lock-in. Interoperability is the ability to use customer systems and management tools to interact with multiple cloud services. Both terms present challenges that cloud governance should address.
- **SLA**—A service level agreement is a contractual representation to define, measure and monitor the qualitative and quantitative characteristics of the cloud service (e.g., performance and security).

When building a governance policy, an organization should consider all relevant aspects.

Tip: The CSA cloud octagon model is a framework for cloud adoption that describes the various roles and processes required for cloud adoption in an organization.

Becker and Bailey's Cloud Governance Dial[1] focuses on the IT governance domains applicable to cloud adoption for an associated product, as follows (**figure 1.6**):

1. Process—What is moving to the cloud

2. Delivery—IT governance domain objectives and deliverables

3. Deployment—How it will be delivered

4. Cloud formation—How it will be deployed

5. ERM—The cloud formation

6. Control—Cloud governance

Figure 1.6—Becker and Bailey's Cloud Governance Dial

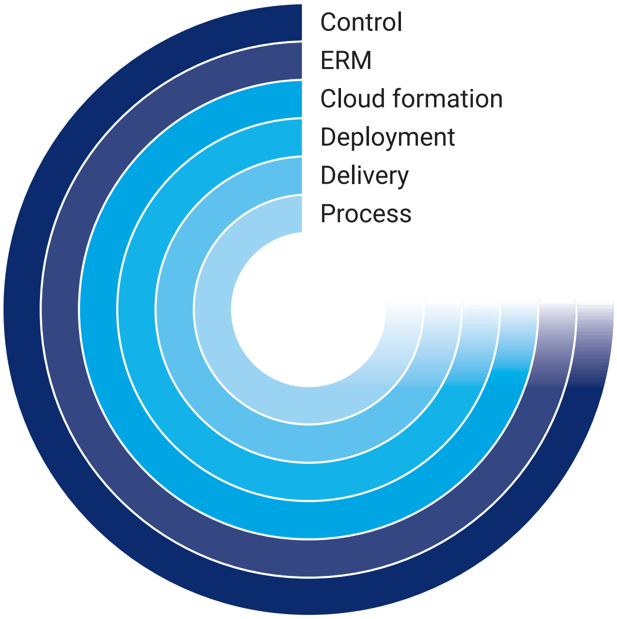

Source: Becker, J.; E. Bailey; "A Comparison of IT Governance & Control Frameworks in Cloud Computing," AMCIS 2014 Proceedings, August 2014, https://pdfs.semanticscholar.org/87f4/71d5e562e7e7640ce97bcf96b57ce0b02a32.pdf

[1] Becker, J.; E. Bailey; "A Comparison of IT Governance & Control Frameworks in Cloud Computing," AMCIS 2014 Proceedings, August 2014, https://pdfs.semanticscholar.org/87f4/71d5e562e7e7640ce97bcf96b57ce0b02a32.pdf

1.3.7 Identify All Governance Stakeholders in an Enterprise and Third-Party Players and Bodies

In addition to the roles identified in section 1.3.5 (i.e., IT, GRC, legal and procurement departments), other internal and external stakeholders play a valuable role in the cloud governance process:

- **Business owners**—Set the business requirements
- **Enterprise architects**—Set the strategy and technical requirements
- **Senior management and board**—Approve governance considerations, hold accountability
- **Security department**—Set security requirements and frameworks
- **Cloud providers**—Manage their end in the shared responsibility model
- **Cloud integration/implementation brokers**—Responsible for building the cloud infrastructure, for cloud migration and, sometimes, for monitoring the post-migration infrastructure
- **Software development contractors**—Responsible for building the application that will run on top of the cloud service
- **Security consultants**—Responsible for setting the security and compliance requirements, approving the solution architecture, and implementing security tools
- **Auditors**—Responsible for the security review of the cloud platform and application, usually against a predefined agreed-on set of controls and organizational requirements

1.3.8 Approaches for Assessing Assurance, Quality and Performance of Cloud Services

When implementing or migrating to a cloud service, customers need assurance that:

- Their data is protected from unauthorized access.
- Their data is collected and processed only according to agreed-on purposes and to applicable laws and regulations.
- The service is being provided according to the specifications defined in the contract.
- Disruptions will not affect their ongoing business.

Customers implement cloud assurance programs to verify these requirements.

One way to describe cloud assurance: the process of providing confidence that a cloud service will perform during its lifecycle according to established policies, specifications, declared features and SLAs, and in compliance with applicable laws and regulations.

Cloud assurance is often based on distinct mechanisms:

1. Contracts and terms of use, including service level agreements
2. Provider self-assessment and technical documents (e.g., specifications) and due diligence questionnaires
3. External attestation and certification audit reports (e.g., SOC2, ISO27001)
4. First-party risk assessment against a set of controls, sometimes with an independent auditor
5. Provider reputation
6. Provider financial stability and market value
7. Provider cyberinsurance

When contemplating the cloud assurance process, customers need to consider the following:

- **Provider maturity and ability to execute**—An important factor in the assessment process. Although maturity is an elusive term, it plays an important role in customers' decisions. Mature providers are likely to be more transparent, more experienced and in a better position to provide assurance. Evaluating provider maturity involves assessing the following:
  - Fiscal performance, market share, company size, number of employees and market positioning
  - Board and senior management involvement, supervision and support
  - Transparency with customers, including financial and technical transparency
  - Number of years in the specific market segment
  - Market vision and development road map
  - Implementation of a formal process to measure service performance improvement
- **Provider security**—Includes every aspect of provider development and ongoing operations, encompassing both back-office security (IT security) and the actual cloud service security program.
- **Provider compliance**—Responsibility for adhering to the requirements of all laws and regulations applicable to the customer. Providers are also responsible for their subcontractors and service providers, and the level of assurance they provide. Further, providers are required to adhere to their own laws and regulations, so customers must verify that there are no conflicts.
- **Provider privacy**—Responsibility to adhere to applicable privacy laws and regulations, managing restrictions on data processing, storage and transfers. Customers need to verify that provider privacy policies align with their own privacy requirements, including data residency obligations.
- **Provider monitoring**—Important for detecting changes in providers policies, governance schemes and ability to execute.

Tip: Different services from the same provider can have different security features, levels of compliance and certifications. The risk management process should evaluate the specific cloud service features and not just the CSP general security posture.

1.3.9 Cloud Assurance Program Development

Cloud assurance program development usually takes place in three phases (**figure 1.7**).

Figure 1.7—Three Phases of the Cloud Assurance Program

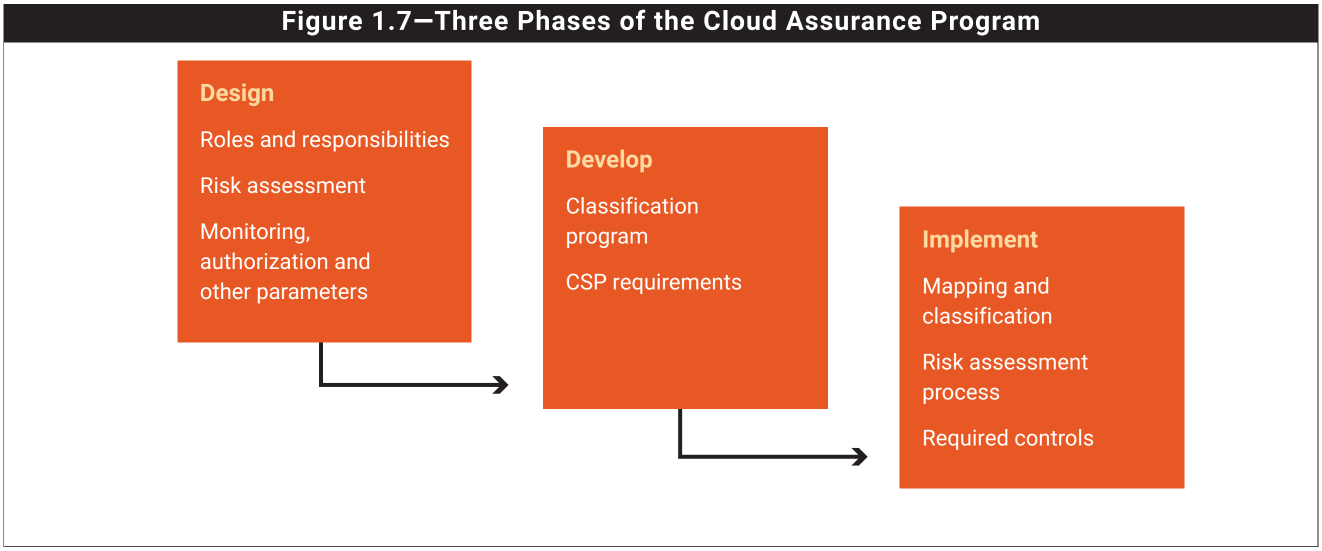

Design

Roles and responsibilities

Risk assessment

Monitoring, authorization and other parameters

Develop

Classification program

CSP requirements

Implement

Mapping and classification

Risk assessment process

Required controls

- **Design an assurance program**
  - Identify roles and responsibilities among the relevant stakeholders, including the oversight role (responsible for success of the assurance program) and the program operations role (responsible for the daily operations of the program).
  - Design the risk assessment program for cloud migration—all deployment models (private vs. public), service models (IaaS/PaaS/SaaS) and data classification models affect the risk management process. Organizations should list the cloud risks they foresee, and then examine the designated cloud service against those risks to determine risk likelihood, impact and tolerance.
  - Design the program authorization process by approving new services; addressing unique risk scenarios; approving architecture, designs, controls and tools; and approving exceptions.
  - Design the monitoring goals and requirements. At this stage, the customer is designing the metrics to measure the success of the cloud deployment and the continuous improvement process.
  - Design the monitoring and feedback process. Monitoring can be periodic, scheduled for a point in time, or continuous and always-on. Deciding on the correct monitoring method takes place during the design stage, but it can be changed later. In addition to continuous monitoring, continuous improvement and feedback are essential for the risk-based approach and for the governance process.

- **Develop the assurance program**
  - Develop a data classification program. Good data classification techniques are fundamental for any cloud governance policy, because the actual process of risk management depends on the classification of the data. (The higher value of the data, the more effort should go into the risk management program.) In the data classification process, consider the data life cycle, including possible changes in data definition. How a certain data set is defined in the current environment may differ from how it might be defined in the future, as a consequence of new laws, regulations, standards or technologies. For example, the definition of consumer broadened under the California Consumer Privacy Act (CCPA), which has an impact on the risk management efforts of organizations affected by the new law.
  - Define requirements for CSP and CSP services. Usually there are sets of different requirements based on the classification of the data. This is also the juncture to elevate the use of international standards and third-party certifications or attestations—e.g., sensitive information will migrate only to providers that obtained third-party evaluation of the effective controls based on SOC2 requirements. Requirements should be made for both the CSP as a company and its relevant cloud services. Different services, either from the same CSP or multiple CSPs, can demonstrate different security characteristics, so it is important to evaluate each provider service separately.
  - Identify the laws, regulations and standards that are relevant to different data types. This process depends on variables, such as data classification (e.g., personal data vs. business data vs. financial data); CSP jurisdiction (e.g., where is the CSP headquartered? Where is the data collected? Where is it processed? Where is it stored?); and cloud customer industry sector and jurisdiction.

- **Implement the assurance program**
  - Map and classify the data relevant for the migration. The mapping should include roles, such as data owners and stewards (responsible for data governance, usually a business role); administrators (responsible for actions such as backup or monitoring, usually a technical role); and user categories (people or units who actually need different degrees of access to the data and services).
  - Identify the processes, functions and systems relevant to the migration. Data is a major part of the cloud migration, but successful cloud migration also involves identification of the systems that will provide access to the data, and processes and functions that enable data updates, backups and access controls.
  - Rate cloud application sensitivity (usually how sensitive the application is to confidentiality and integrity threats) and criticality (usually the importance of application availability). GDPR added resilience to the CIA triad, so rating the cloud application ability to function under adverse conditions is part of the resilience factor. Based on the organization cloud strategy, choose the cloud delivery service level (IaaS/PaaS/SaaS) and deployment models (private/public/hybrid/community). Some applications can be accessed as SaaS, and some need to be developed on top of IaaS/PaaS. Some workloads can use public clouds, while some might stay in the private or hybrid zone.
  - Implement security requirements and risk assessment processes. Following analyses of application sensitivity, technical characteristics, and relevant laws and regulations, organizations can start implementing security architecture, mitigating controls and specific requirements.

Example: The UK government adopted a new policy about the security classification of information assets. The new policy narrowed the classification levels from six categories to three (official, secret and top secret). With the implementation of the new classification scheme, the services provided through the G-Cloud framework will fall into the official category.[2]

1.3.10 Tools and Techniques to Design, Implement and Operate a Governance Program

To provide assurance, various tools and techniques are available. The following is a list of tools for use in cloud governance:

- **Contract**—Providers usually require customers to sign a customer agreement before using services. Those agreements are a major foundation of governance and assurance, because they provide controls on the relationship with the CSP. Cloud agreements or contracts usually consist of service terms (SLA, acceptable use policy, technical support) and legal terms (jurisdiction, dispute handling, remedies). Those terms are the foundation of the shared responsibility model, which is usually not described directly in the contract. In general, cloud providers rely on unified service and do not negotiate their contracts for every customer request. However, large organizations probably will be able to get more changes, and smaller providers probably will demonstrate more flexibility. See section 1.4.13 for more information about contracts as a governance tool.
- **Security assessment**—Assessment of the cloud service and provider security posture is critical when ensuring governance. There are many assessment triggers; it can be a preliminary process before the actual cloud migration or an ongoing process to monitor cloud providers. It can result from corporate policy or from regulatory guidelines. Security assessments range from reviewing papers to actual hands-on red teaming

[2] European Union Agency for Cybersecurity, "Security Framework for Governmental Clouds, Annex A & B Case Studies and Interviews," www.enisa.europa.eu/publications/security-framework-for-govenmental-clouds/annex-a-b-case-studies-and-interviews/view

penetration testing on the overhaul platform. In general, the goals of each assessment process are to understand the relevant controls; to check that customer-side controls are in place and operating correctly; and to verify that the customer is asking the cloud provider the right questions about responsibilities under the shared responsibility model. See section 1.4.14 for more information about security assessments as a governance tool.

- **Audit overview**—The difference between assessments and audits is that audits usually reflect a much more in-depth process. Audits are based on predefined scopes and focus on key controls that are usually associated with specific risks. Audits require a lot of planning and preparation. At the end of the audit process, findings should be presented and prioritized, based on risks. See section 1.4.15 for more information about audits as a governance tool.

Relationship of Governance Tools to Risk

The question of which governance tool to use depends on the following factors:

- **Service model**—When selecting SaaS, the contract is almost the only way to ensure the existence of specific controls. When using PaaS or IaaS, most controls are the responsibility of the customer, so using audit overview on the implementation might produce better results.
- **Identity of the provider and the service**—Some providers demonstrate better security maturity than others. Customers may decide to rely on contract terms only with providers that have rigid security policies and can demonstrate higher maturity, but they may choose to examine other providers through a thorough audit process.
- **Risk profile**—The higher the risk to the customer workloads, the more effort goes into the governance process. When dealing with high-risk workloads, customers might request controls in the contract, carry out a security assessment and continue with ongoing audits.

1.4 Cloud Trust, Transparency and Assurance

Trust in the cloud is a function of transparency, assurance and accountability This section discusses these variables.

1.4.1 Defining Trust

Trust is often used as a general term covering security and privacy from the provider.

Jingwei Huang and David M. Nicol define trust as follows:

> *Trust is a mental state comprising: (1) expectancy—the trustor expects a specific behavior from the trustee (such as providing valid information or effectively performing cooperative actions); (2) belief—the trustor believes that the expected behavior occurs, based on the evidence of the trustee's competence, integrity and goodwill; (3) willingness to take risk—the trustor is willing to take risk for that belief.*[3]

Huang and Nicol identify two types of trust, based on the trustor's expectancy: "Trust in performance is trust about what the trustee performs, whereas trust in belief is trust about what the trustee believes. The trustee's performance could be the truth of what the trustee says or the successfulness of what the trustee does."[4]

Categories of Trust

Trust in cloud computing is divided into the following categories (**figure 1.8**):

- **Reputation-based trust**—The reputation of an entity is the overall estimation of the public's trust toward it. Generally, many entities in a community trust an entity that has a high reputation. An entity that is required to

[3] Huang, J.; D. Nicol; "Trust mechanisms for cloud computing," *Journal of Cloud Computing*, 24 April 2013, https://journalofcloudcomputing.springeropen.com/articles/10.1186/2192-113X-2-9

[4] *Ibid.*

build trust decisions on a trustee uses the reputation to compute or approximate the trust level of the trustee. The reputation of cloud affects the selection process of cloud services. Therefore, CSPs aspire to build and preserve a higher reputation. Reputation is classically represented by a broad score reflecting the overall outlook, or a small number of scores on numerous foremost aspects of performance.

- **SLA verification-based trust**—In SLA verification-based trust, after establishing preliminary trust and accessing a cloud service, the cloud customer is required to validate and re-examine the trust value. A service level agreement is a lawful contract between the two communicating parties: customer and provider. Therefore, monitoring the quality of service (QoS) parameters and verifying SLAs are essential steps in trust management for cloud computing. A third CSP party is required to provide these types of services.
- **Policy-based trust**—Policy-based trust requires a formal trust construct. In a related area, public key infrastructure (PKI) is a widely used technology that employs formal trust methodologies to support key certification, digital signing and validation. It also supports data attribute certification and validation. In this category, the trust in a certification authority (CA) is dependent on its confirmation with defined policies governing delivery, retention and validation of public key certificates. Certificate policies play a key role in PKI trust.
- **Evidence-based trust**—In evidence-based trust, a truster's belief in a trustee's predictable behavior is based on proof supporting attributes of adeptness, helpfulness and honesty.
- **Societal trust**—Societal trust refers to any individual or organization. In the cloud, each entity must be trusted.[5]

Figure 1.8—Categories of Trust in Cloud Computing

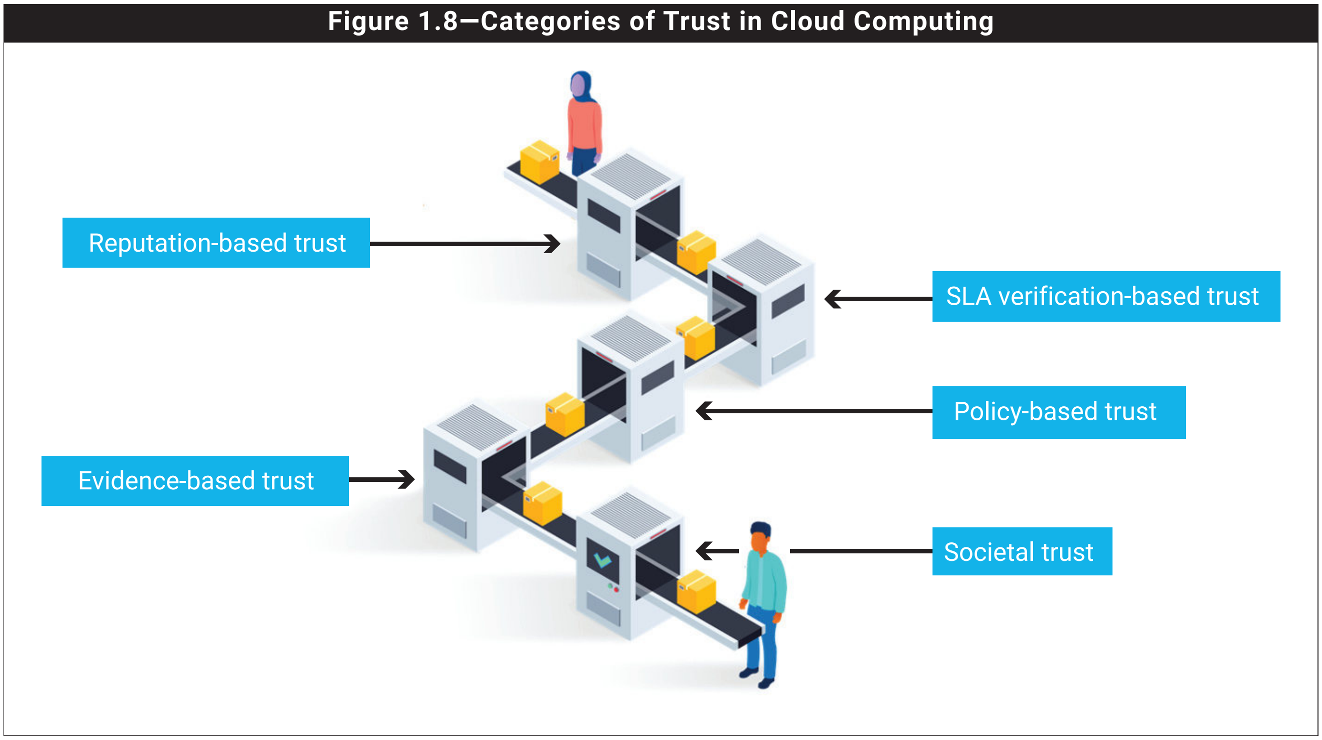

To increase trust between all relevant parties to the cloud migration, the cloud governance team must first identify and address the material needs of the stakeholders:

- **Consumer-side business, technology, risk and control functions**—Need a way to understand and make choices about how data may be used in the cloud, and the consequences and risks associated with those choices.

[5] Ritu; S. Randhawab; J. Sushma; "Trust Models in Cloud Computing: A Review," *International Journal of Wireless and Microwave Technologies*, April 2017, www.researchgate.net/publication/318583494_Trust_Models_in_Cloud_Computing_A_Review

- **Providers, integrators and brokers**—Must ensure that data is treated appropriately, and that they retain control over how it is used.
- **Regulators**—Must understand the legal frameworks that apply, and the ways to hold providers and customers accountable for what happens to data.

1.4.2 How to Define Trust in the Cloud

Trust in the cloud refers to the superset of beliefs and expectations of the three primary stakeholders across the five trust categories previously outlined. The trust conferred by an individual organization customer to a specific CSP is based on a range of trust category factors that include its evolving perception of the brand, the level of open source evidence available—media reports, feedback from other customers of the provider, or information from peers in the industry and other practitioners—assertions the CSP makes, and a focused examination of the evidence it offers to back up those assertions. Going further, and assuming an organization customer selects a given CSP, it may subsequently limit its approval to a subset of the CSP services, until certain road map items and assurances are achieved.

Therefore, cloud governance should facilitate and support the internal stakeholders' journey to conform their beliefs and expectations (which may be based on anecdotal, selective or biased information) to a more rigorous and formalized process. From this foundation, decision makers can credibly define the levels of trust for different potential use cases, supported by clear transparency and control requirements. The supporting process should result in broadly similar outcomes for the same inputs, so that another organization could follow the same threads of logic and arrive at a broadly similar conclusion, given the same definition of risk appetite, control requirements and service levels required.

Trust can be present when independent parties are engaged to inspect chosen components or services and, based on their assurance work, make statements that reflect the truth of the CSP assertions.

Trust works both ways: An organization may need to prove a level of trustworthiness before a CSP will share certain information. Assurance, transparency and accountability apply equally to both parties.

A critical mistake or misplacement of trust can occur when an organization believes its CSP is responsible for all security. In most cases, that is not the situation—and in fact the CSP has documented the need for the organization customer to carry out certain functions and apply certain controls. Ideally, the customer should have a document or RACI model that maps out internal roles and responsibilities. Management needs to clearly demonstrate that there are no gaps between the organization's responsibilities and those of the CSP.

The shared responsibility model increases the complexity of the relationship between various actors in the supply chain. For example, a customer is faced with a supply chain comprising three actors—a SaaS provider building on top of an IaaS provider, plus a third party that has built a plugin or extension to interface with the SaaS (**figure 1.9**). The challenge then becomes about establishing multiparty trust.

Figure 1.9—Security Posture and Shared Responsibility in the Cloud Supply Chain

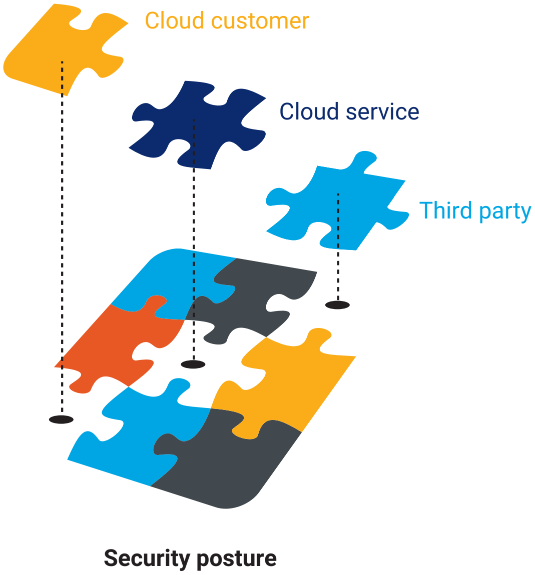

1.4.3 Ways to Measure Cloud Trust

There are quantitative and qualitative ways to measure cloud trust.

- **Desktop review**—Examines the documentation the CSP makes available, including security statements, descriptions of controls and assertions by independent third parties that had sufficient access to validate their assertions.
- **Experiential measurement**—Derives from the customer's own experience of the service, e.g., having engaged in a sandbox exercise to experiment with a cloud service. The organization then tries to link the description of the services and the CSP promises with the reality of integrating and testing some of those services. An organization with internal technology experts can write gap lists and shortcomings. In regulated industries, certain circumstances may warrant a direct peer relationship between the cloud customer and the CSP internal experts. This can be a function of the size of the organization customer—the larger it is, the deeper the level of peer relationship needs to be to build and measure the high levels of trust necessary for that organization to hand more workloads to the CSP.
- **Secondary sources**—Relies on professional associations or peers in an industry sector. Often, the source is through consultants or hires who bring new skills in-house and have experience and knowledge gained from prior work with a CSP. Industry analyst reports can help to drive awareness of a CSP beyond trust, but there may be specialist reports analyzing a CSP from technical/procedural/performance perspective.
- **Reported incidents**—Influence trust informally, e.g., if a breach occurs and becomes public, the incident itself and the CSP response may affect consumer trust.
- **Supply chain intelligence providers**—Use new ways to respond to questions, enabling customers to achieve a level of understanding and assurance that is foundational for developing trust in a service provider. There are now third-party intelligence services that monitor, score, risk-rate and report on service providers, including CSPs. Some have a service that gathers CSP responses to commonly submitted customer questions into a portal, where prospective customers can review the CSP existing answers and then submit additional questions specific to their needs. These brokers may also carry out operational monitoring of a service for malware reports and

provide a score—or if there has been a breach, they might submit Freedom of Information Act (FOIA) requests to try to score how well a breach was handled.

- **CSA STAR**—The CSA Security Trust Assurance and Risk (STAR) initiative enables CSPs to validate their cloud security while offering proof of the controls they have in place to current and future customers. STAR helps customers assess which organizations meet the level of assurance they require.

Tip: Some of these measurement methods scale well—others do not. Another point to be aware of is the level of commoditized trust indicators compared to very niche assurance and trust conversations. This is a spectrum. It is also worth noting that highly involved engagements between a CSP and a customer usually carry a high cost.

1.4.4 Impact of the Cloud Shared Responsibility Model on Trust

The shared responsibility model is key to the cloud. It imposes clear responsibilities on the various parties involved, but it is not up to any party to evaluate if others are keeping to their side of the bargain (section 1.2.5 elaborates further on roles and responsibilities). In terms of transparency within their own organization, business owners need to have a clear line of sight on how the internal IT team has mapped the shared responsibility model across the different control functions within the organization.

An organization should define the system boundaries by identifying resources and security controls required for all the cloud services. The compilation of a RACI matrix, establishment of clear accountability and documented procedures to handle well-understood events will provide clear sight and reduce the impact of the shared responsibility model. Leaders can trust that their own people are doing what they are meant to do by having them produce evidence and put measures in place, providing ongoing assurance that the right people are doing the right things at the right time. If any of those controls are not working fully, there needs to be accountability.

The organization needs to carry out an appropriate level of discovery for the number of workloads it plans to deploy or services it will consume in the cloud—e.g., how many people are needed to operate the controls so the service is trustworthy? Certain decisions flow from this: Does the organization need to scale up internally by hiring full-time staff or working with external consultants? Within an organization, the responsibilities need to be explicitly stated and reflected in the way the business runs its own operations, to avoid gaps between where its responsibilities end and the CSP responsibilities begin. It may be useful to outline those responsibilities—and their exact interpretations—in an agreement between both sides.

Cloud services change much more rapidly than traditional IT services, with new services released and old services deprecated. The shared responsibility model can evolve over time as the provider adds new services and upgrades existing services. For example, if a CSP decommissions a particular service offering, it is important to understand who is responsible for managing the lifecycle of data migration and conversion to the new service. All parties need to consider their respective responsibilities through the life cycle of services, not just at a point in time.

A cloud customer may be increasing its risk based on the nature of the cloud service. Noncore services in the cloud can be at risk of being replaced or substantially changed. The dependency chain with core services can give confidence that they will continue for some time to come. Changes to fundamental services carry a much higher cost to all users.

1.4.5 Relationships Among Transparency, Assurance and Accountability

Trust in the cloud is a function of three variables: transparency, assurance and accountability. It may be helpful to think of these elements as the three legs of a stool. If there is strong transparency but limited accountability and a level of moderate assurance, that will affect how much the customer trusts a CSP. Those three elements are, effectively, in the eye of the beholder. They also work both ways: The organization must know its own risk appetite—whether it is subject to any regulations, and what it is prepared to accept—because it is unlikely any CSP can fully deliver on all three criteria. An organization that is subject to contractual or regulatory requirements may not be able to use a service for certain workloads if one of those three elements is weak (**figure 1.10**).

Figure 1.10—Relationships Among Cloud Transparency, Assurance, Accountability and Trust

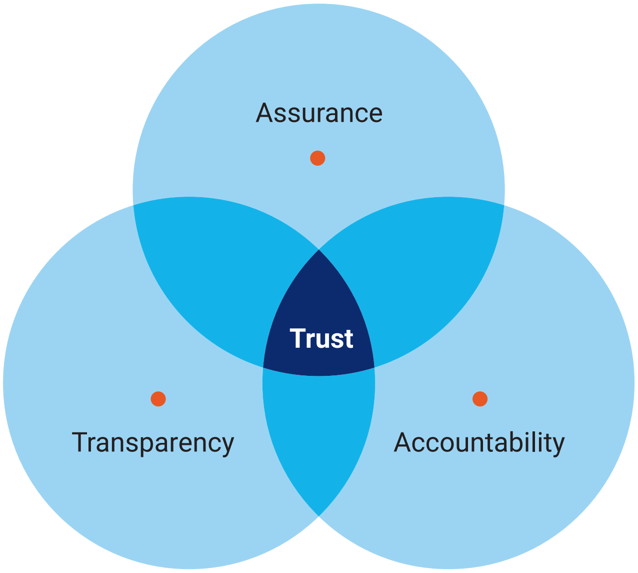

Following are descriptions of the three variable of trust:

- **Transparency**—In the absence of providing trusted access to its platform, a CSP can show evidence at a level that permits cloud customers to make reasonable judgments about what the provider does and does not do. The more factors a CSP can demonstrate—e.g., restricted access, encryption or other controls—the more convincing that it operates transparently. For larger accounts, the level of transparency can evolve during the relationship, as the CSP gains a deeper appreciation of the business possibilities.
- **Assurance**—The knowledge that a control is operating as stated can come with different levels of assurance, from a design review to an operational check, to a deep audit and ongoing inspection. It may be linked to frequency—e.g., how often are the controls checked? Is monitoring continuous or does it occur at a point in time? It can be linked to who performs the check—e.g., is it the CSP, the cloud provider assessor, or a totally independent auditor? Assurance is covered in depth in chapter 6.

Tip: Some CSPs have monetized assurance by making it a service offering. By subscribing to optional or enhanced programs, the CSP customers get a level of access to security information that is not publicly available to all. This helps them to assign a specific budget number to security. Even customers that do not subscribe may benefit from this rising tide, because the CSP may provide additional security to core services regardless of whether the user pays extra for it.

- **Accountability**—Is the CSP demonstrating a commitment not only to state that it does something, but also to show that it does something? If a service is not operating as stated, accountability is reflected in how the CSP adapts and responds to make amends. Accountability also relates to the risk management and remediation processes that should ensure identified risk is managed appropriately and remediated in a timely manner.

1.4.6 Transparency

Transparency has considerable importance and attention due to a lack of visibility in cloud services, which is a business risk in terms of security, compliance and governance. Private organizations and government institutions gain the confidence of the public and customers by being more transparent about their motives, goals and processes for accessing data.

Transparency is hard to define because it means different things in different scenarios:

- In fiscal economic terms, transparency is defined as governments being open with the public about structures and functions, policy intentions, public sector accounts and projections.
- In social sciences, transparency connotes the ability of interested parties to see otherwise private information.
- In information technology, transparency refers to actions aimed at improving openness and accountability.[6]

This section addresses transparency as the practice of disclosing information related to security practices and controls used to protect customer data and applications hosted in the cloud environment. It entails the provision of synchronous and asynchronous information about events related to data and applications to give sufficient visibility into incidents affecting their assets while in the custody of a CSP.

Transparency is a major foundation for creating trust. It is very difficult for a cloud customer to trust cloud providers without full transparency into the provider's security posture, security considerations and technical capabilities.

Categories of transparency include:

- **Proactive transparency (voluntary)**—The CSP voluntarily discloses information to users by means of autonomous agents or manual processes. It stems from CSP initiatives to make security information about their offerings available to all without requiring that a request be filed. It can be generated through various methods, such as notifications or reports, including content on websites and portals, benchmarks, and white papers. This disclosure is generally intended to improve customer trust and confidence, to help customers evaluate the basic control environments of a CSP, to demonstrate its certification to meet regulatory or industry requirements, and to help customers address specific questions around general practices. It is not meant to reveal information that could jeopardize a CSP security posture or expose it to harm. For instance, when CSC data fall under restrictions emanating from regulatory or compliance requirements, the choice of a CSP hinges on being satisfied that it is fully compliant. Otherwise, there is the risk of violating regulatory, legal or other privacy requirements. One of the benefits that accrues from a CSP initiative to proactively disclose information is that all CSCs, including small, large and medium organizations, get timely access to details without needing to file special requests, which can involve significant commitments.
- **Reactive transparency (necessary)**—A CSP is expected to respond to a user's specific request. It emerges from a request-response routine in which the CSP must provide additional information to the CSC. Through the request-response regime, a prospective user files a request and receives information from an existing CSP or asks for an incident notification to be sent to the CSC. Generally, the contents of a CSP response must be more than a meager attempt to beat competition from other vendors. They must reflect the actual settings, offerings or security statuses of the CSC assets—information that indeed increases transparency about how it distinctly addresses individual requirements. The advantage of the request-response approach is its ability to enable users to specify their exclusive security requirements and for the CSP to identify suitable controls. However, it arguably has become a traditional practice for CSPs to frequently publish information in the public domain so that future cloud users do not have to file a request, saving time for both the CSP and customers.
- **Contractual transparency (statutory)**—A CSP is often mandated by law to provide transparency while providing services to CSCs. A valid written agreement may require the CSP to make full disclosure of all

[6] Ismail, U.; S. Islam; M. Ouedraogo; E. Weippi; "A Framework for Security Transparency in Cloud Computing," Future Internet, February 2016, www.researchgate.net/publication/295082455_A_Framework_for_Security_Transparency_in_Cloud_Computing

essential security services on an individual basis, while refraining from divulging information that could compromise other users' privacy. SLAs are tools that CSPs and CSCs use widely to ensure transparency. They establish a pact for aligning the security requirements requested by a CSC with the security levels offered by a CSP. The SLA forms the basis for defining responsibilities and the remedies available to customers in case of a contract breach. Fundamental aspects of the SLA include representation of the contexts shared by the two parties, and how each actor utilizes the contexts in its own operations throughout the SLA. The SLA provides comprehensive description and transparent security processes for both the CSP and customer to avoid uncertainty, apprehension and disputes. Conventional SLAs generally provide clarity on CSP service offers, unambiguous definitions of expectations and obligations on both sides, and the boundaries of liability. A notable limitation to this class of transparency is that essential security properties of a customer may not be captured.

- **Cloud security transparency deployment practices**—There are different dimensions to information disclosed through transparency. Some information disclosure supports the CSC in decision making. In some cases, certain disclosures present nonessential information that may create ambivalence. Disclosures fall into two categories: opaque and explicit transparency.[7]

Most CSPs today invest heavily in creating transparency through use of the following tools:

- **Security policies**—A rigid, transparent, wide and honest security policy is mandatory for CSPs. Security policies are usually built from a collection of different policies: incident management, change control and patch management, disaster recovery, and so on. Some CSPs may stipulate that the customer must sign a nondisclosure agreement in order to get access to more sensitive parts of those policies.
- **Service level agreement**—The CSP SLA is an important tool in understanding the availability risks (from the confidentiality, integrity and availability triad) to the application and data.
- **Self-assessment, third-party assessment and certification**—Many CSPs use a combination of assessments, audits and certification to provide reliable details about their own security postures and risk management processes.
- **Right to audit**—Some providers allow customers to perform their own audits, tests and scans. This varies from one provider to another, and types of audit can range from limited remote vulnerability scans to comprehensive walk-in audits on site. When customers are allowed to exercise the right to audit is also important.

Tip: Trust and transparency depend heavily on the detail and accuracy of the information provided by the CSP and its auditors. Experienced customers are likely to reject security policies that pretend the CSP services are flawless and do not address the continuous improvements that any organization must make.

1.4.7 Accountability

> *Accountability can be formally defined as being the state of accepting allocated responsibilities, explaining and demonstrating compliance to stakeholders and remedying any failure to act properly. Responsibilities may be derived from law, social norms, agreements, organizational values and ethical obligations. Accountability is not an absolute; it is only meaningful in the context to which it is applied (i.e. what one is accountable for). Accountability is then associated with specific principles and actions corresponding to that context.*[8]

According to the Cloud Accountability Project (A4Cloud), the following qualities are the foundations of accountability:

- **Transparency**—The property of a system, organization or individual of providing visibility of its governing norms, behavior and compliance of behavior to the norms

[7] *Ibid.*
[8] Cloud Accountability Project (A4Cloud), "Aspects of Accountability," www.cloudaccountability.eu/content/aspects-accountability

- **Responsiveness**—The property of a system, organization or individual to consider input from external stakeholders and respond to their queries
- **Responsibility**—The property of an organization or individual, in connection with an object, process or system, of being expected to take action to comply with norms
- **Remediability**—The property of a system, organization or individual to take corrective action and provide a remedy for any party harmed in case of failure to comply with governing norms[9]

Building cloud accountability is a challenging task because cloud computing involves many stakeholders from inside and outside an organization, and cloud characteristics such as cross-border data transfers and shared responsibility add complexity. Tools such as a RACI table can assist in clarifying accountability. See section 1.3.5 for more information about RACI tables.

1.4.8 Cloud Governance Framework

The requirement to build a cloud governance framework and not reuse an existing IT governance framework comes from the fact that cloud is a new technology model with unique characteristics. This section discusses the cloud governance framework.

Risk Specific to Cloud Operations

Among those unique characteristics are increased reliance on third parties, unique automation features, shared resources, data crossing between jurisdictions, self-provisioning of resources, highly connected platforms and increased use of virtualization.

Those characteristics produce new risk, threats and attack vectors. Following are some examples:

- **Provider administration**—CSPs are responsible for securing large parts of the infrastructure and the application. CSPs failing to perform their responsibilities present significant risk.
- **Virtualization risk**—Cloud services rely heavily on virtualization. Some virtualization platforms used are mature and well-known (e.g., operating system virtualization), but some are new and less mature from a security aspect (e.g., containers). Attackers keep exploring ways to manipulate virtualization environments to detect vulnerabilities or misconfigurations.
- **Compliance risk**—A complex supply chain, data managed by three parties, data crossing jurisdictions and shadow IT are cloud characteristics that can lead to breaches of compliance efforts under certain circumstances.

See section 1.4.10 for more information on cloud risk and risk management methodologies.

The previously mentioned risk may not be exclusive to cloud computing, but it is very common to it, and it may force cloud customers to deploy a new IT governance model.

Cloud-specific Security Controls Framework

Cloud computing tends to produce sets of controls that differ from the regular controls that organizations are familiar with—from hosting or outsourcing, for example. Because cloud computing delivers new business models wrapped in new technologies, it is necessary to implement several types of controls—administrative (contract, SLA), technical (multifactor authentication (MFA), encryption) and physical (secure data center).

A cloud-specific security framework incorporates key controls, including some that are unique to cloud platforms, such as virtualization security and automation.

[9] *Ibid.*

A cloud-specific controls framework is recommended over traditional security frameworks that need to be mapped to cloud controls, which entails an additional adoption process and the risk of leaving out elements.

Cloud-specific security controls frameworks include the following:

- CSA Cloud Control Matrix (CCM)
- ISO/IEC 27017:2015
- BSI C5
- NIST 800-53
- PCI SSC Cloud Computing Guidelines

See section 2.6.2 for additional details about these frameworks. Chapter 3 is dedicated to CCM.

Shared Responsibility Model Within the Cloud Governance Framework

Cloud security is based on the shared responsibility model, which describes the sharing of responsibilities between the cloud provider and the cloud customer. Some responsibilities are placed on the provider and some on the customer. The sharing characteristics differ for IaaS, PaaS and SaaS (**figure 1.11**).

Examples of controls that are usually the responsibility of the cloud provider follow:

- Physical security of the data center
- Security of the hosts and hypervisor layers
- Background checks of provider employees
- Access of providers employees to customer data

Examples of controls that are usually the cloud customer responsibility follow:

- Customer data
- Interfaces and management plane security (MFA, access controls)
- External cloud backups

Figure 1.11—Cloud Customer and Provider Shared Responsibility Risk Matrix

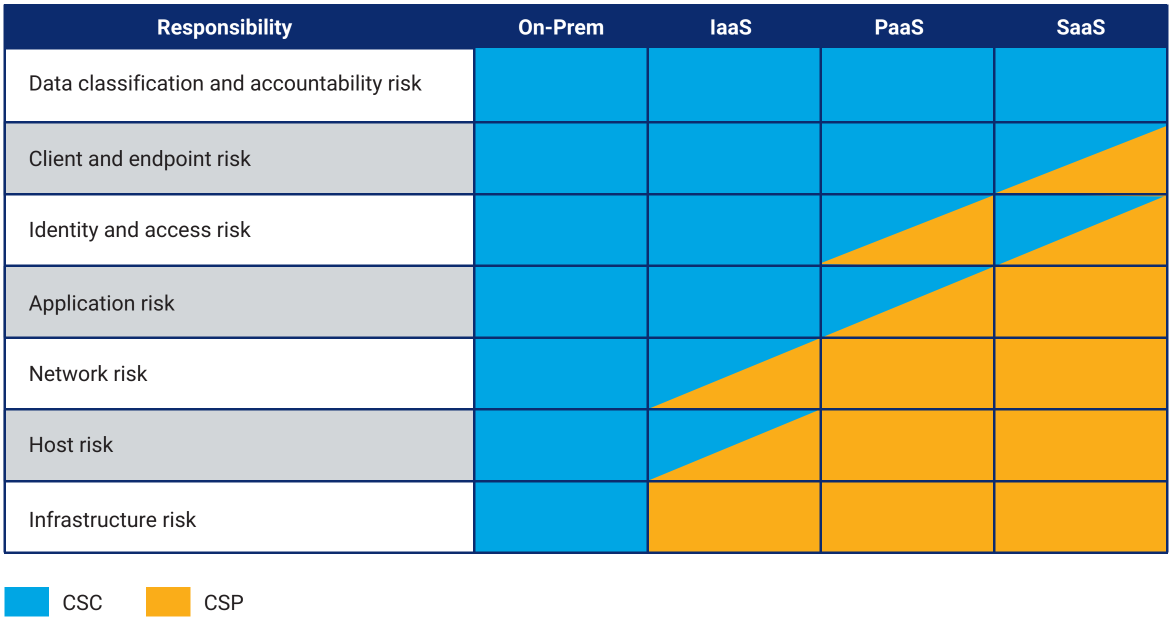

Responsibility	On-Prem	IaaS	PaaS	SaaS
Data classification and accountability risk	🟦	🟦	🟦	🟦
Client and endpoint risk	🟦	🟦	🟦	🟦/🟧
Identity and access risk	🟦	🟦	🟦/🟧	🟦/🟧
Application risk	🟦	🟦	🟦/🟧	🟧
Network risk	🟦	🟦/🟧	🟧	🟧
Host risk	🟦	🟦/🟧	🟧	🟧
Infrastructure risk	🟦	🟧	🟧	🟧

🟦 CSC 🟧 CSP

Source: Microsoft TechNet, "Shared Responsibilities for Cloud Computing," 25 October 2019, https://gallery.technet.microsoft.com/Shared-Responsibilities-81d0ff91

The cloud governance framework requires an understanding of the shared responsibility model and how to implement it correctly. When analyzing the shared responsibility model, organizations usually need to document the following:

- What are the relevant controls to be implemented? Organizations usually adopt a controls framework, which can be non-cloud specific (e.g., ISO27001) or cloud specific (e.g., CSA CCM).
- Who is responsible for implementing each control—the cloud provider, the customer or perhaps a third party? Often, the answer is that both the customer and provider are responsible for a specific control. Such cases need a more detailed analysis that breaks the control into lower-level controls.
- Who is responsible for each cloud customer control within the organizations, how the control is monitored, and which key performance indicators (KPIs) exist to test operational effectiveness.
- Evidence that the provider is doing its part for each cloud provider control, and a means for customer verification.
- Analysis of gaps between provider responsibilities and customer responsibilities and plans to narrow those gaps.

Examples of KPIs to check operational effectiveness of customer controls:

- Number of vulnerabilities remediated on servers (for IaaS/PaaS) in the SLA time frame
- Decreasing number of users without MFA
- Number of servers with data-at-rest encryption

Examples of ways to verify the operational effectiveness of provider controls include:

- **Evidence**—For example, cloud customer receiving the results of a vulnerability scan as evidence of performance
- **Contract**—The provider guarantees in a contract that it will undertake weekly patch management
- **Third party**—An independent external party audit provides assurance of the existence of controls

Governance Framework Driven by Trust, Assurance, Transparency and Accountability

When developing cloud governance programs, organizations must rely on four foundational pillars: trust, assurance, transparency and accountability. The four concepts are indissolubly entangled; the CSA defines trust as a function of assurance, transparency and accountability (see section 1.4.5).

Trust is a critical factor in cloud computing. Without trust, it is very difficult to gain new customers and maintain the satisfaction of existing customers. Yet trust is elusive and hard to measure.

The cloud brought with it the concept of giving up direct control over IT service and security capabilities. While using cloud services, especially public clouds, users must cede direct control over a number of IT, security and privacy-related features to the CSP.

Many security services and capabilities—such as security scanning, penetration testing and audits—might be prohibited by contract. Business continuity and disaster recovery testing are often not permitted. Customer security and privacy policies, including standards and procedures such as hardening, may conflict with the service provider policies. Access to and visibility into key logs and alerts may not be provided.

With the cloud, there is no more direct access to and testing of the physical infrastructure. There is no more direct control over the outsourcer. Instead, there is a new form of governance based on indirect flow of information provided through SLAs, third-party audits and reports on compliance with stringent laws and regulations.

There is a perception that these issues may affect IT governance and governance in general. However, it is more accurate to observe that they have changed the governance paradigm by introducing a new governance model for cloud service providers, based on three key pillars of cloud trust: assurance, transparency and accountability. The challenge to the effectiveness of this new model is the ability to mitigate the unique risk introduced by cloud computing.

To counteract the risks generated by loss of control, governance and risk management approaches should focus on the sources of information used to analyze and assess the cybersecurity risks. Additional emphasis should be given to the variety and accuracy of SLAs, the evidence supporting results of third-party audits and certifications, the quality and details of logging services, etc. It is of paramount importance to reinforce the role of the internal audit function by establishing an auditor mindset and an accountability culture within the organization.

For a governance framework to be effective and efficient, the following should be in place:

- **Assurance**—The customer needs confidence that the provider will deliver the functionality promised and offer levels of service that are measurable, based on clearly defined metrics.
- **Transparency**—The customer needs reliable information to inform or validate its risk decisions. In the cloud, responsibility for some actions (e.g., implementation of security controls) can be transferred to the CSP; however, the cloud customer is ultimately accountable. Transparency is therefore key to exercising due diligence. The principle of trust but verify applies.
- **Accountability**—The customer needs to be able to hold the CSP accountable for its contractual agreements and, when possible, liable for any breaches.

Designing the Governance Framework for Testing the Effectiveness of Cloud Governance Policies

One of the most important governance tools is the ability to test the effectiveness of the existing policy, set benchmarks and KPIs for ongoing monitoring, and plan for continuous improvements.

A good way to start the cloud governance journey is to begin with a minimal cloud adoption policy and test the first applications migrating to the cloud. With maturity and experience, cloud governance evolves as the organization migrates more services.

The initial challenge is to detect the areas where governance can help improve the ongoing process. Different organizations need to govern different aspects of operations. For example, organizations with very dynamic applications might put more effort into cost management, while highly regulated organizations may invest more in security governance.

The following subsections provide examples of issues cloud governance should address, along with some examples of KPIs that are relevant for testing effectiveness.

Cost Management

The complex nature of cloud platform pricing plans, especially in IaaS/PaaS platforms, presents a challenge for ongoing organization operations. Underestimating budgets for new projects, failing to meet budget constraints for existing projects, or making incorrect fund allocations for nonessential or unnecessary cloud services are common pitfalls for cloud adopters. Governance processes should minimize the possibility of those mistakes. Cost management examples follow:

- **Policies**—Build annual and periodic budget forecasts for each cloud project. Perform periodic cost optimization projects to check the largest expenses and plan how to reduce them. Tag resources according to budget relevance. Be aware of providers introducing new pricing plans and the effect on expenses. Investigate budget breaches and one-time losses.
- **KPIs to test effectiveness**—Examine forecast vs. actual budget gaps, annual cloud spending per project, number of operational mistakes, and resulting monetary losses.

Security Policy

Building security governance processes, adding KPIs, and evaluating and monitoring the effectiveness of governance and compliance processes are among the biggest challenges in cloud computing and the foundations of this certification. The lack of control over infrastructure, along with the dynamic and ever-evolving nature of the cloud, make this an ongoing effort rather than a one-time policy creation. Security examples follow:

- **Policies**—Security controls checklist, monitoring policy, incident response procedures and patching policy
- **KPIs to test effectiveness**—Number of incidents, number of exposed servers, number of critical vulnerabilities found, average time for installing a patch

Cloud Agility and Aligning to Business Goals

Organizations move to the cloud for multiple reasons. Some need to reduce costs or risk, some are looking for agile development, some are looking for rapid deployment, and some are looking to increase resilience or adopt continuous improvement methodologies, such as DevOps. The common denominator is that all organizations are turning to the cloud for agility and to support business needs. Therefore, they need to build a governance process to verify that the current cloud adoption program supports the goals they have set. Cloud agility and aligning to business goals examples follow:

- **Policies**—Automated testing, continuous improvement and cloud onboarding policy
- **KPIs to test effectiveness**—Number of cloud applications registered, average time to release a new version and average time for employee onboarding

1.4.9 Cloud Governance Requirements

This section discusses cloud governance requirements.

Scope of Enterprise Requirements

When designing a cloud governance framework, it is important to set the scope of the organization requirements and plan how to implement them. Because cloud computing influences different stakeholders inside the organization and affects many nontechnical aspects, it is important to set the scope so that no area remains ungoverned.

A common mistake is to assume that if an organization is adopting a cloud application, the governance scope covers only the cloud application itself. Cloud providers can rely on other providers and subcontractors, so cloud governance should account for those as well.

Cloud applications can interact with on-premises components and affect different stakeholders in the organization, so cloud governance is always part of a larger governance framework. This is reflected in the scoping process.

In general, cloud governance scope stretches from strategic aspects (i.e., aligning cloud adoption with business goals) to tactics of successful deployments (i.e., provider evaluation criteria) and all the way to details of the ongoing operations (i.e., which control framework to employ).

Tip: Many organizations know they need appropriate SLAs for availability of cloud applications, but they neglect to put the same requirements on their Internet connectivity. This can result in highly available cloud applications that the organization cannot use because of communication line failure.

Different governance frameworks deploy different scoping, with the focus on a variety of aspects of cloud adoption. **Figure 1.12** offers some examples.

Figure 1.12—Comparison of Cloud Governance Models

Standard	Scope	Benefits of Using This Framework for Cloud Computing
COBIT	• Strategic • Environmental • Market • Credit • Operational • Compliance	• Provides a good return on IT-enabled business investments • Manages IT-related business risk • Establishes service continuity and availability • Creates agility in responding to changing business requirements • Achieves cost optimization of service delivery • Lowers process costs • Manages business change • Manages product and business innovation

Figure 1.12—Comparison of Cloud Governance Models *(cont.)*

Standard	Scope	Benefits of Using This Framework for Cloud Computing
ISO27001	Helps organizations comply with numerous regulatory and legal requirements that relate to the security of information	● Establishes the risk management context ● Quantitatively or qualitatively assesses (i.e., identifies, analyzes and evaluates) relevant information risk, taking into account information assets, threats, existing controls and vulnerabilities to determine the likelihood of incidents or incident scenarios, and associated predicted business consequences, to determine risk levels ● Treats, avoids and/or shares risk appropriately, using risk levels to prioritize it ● Keeps stakeholders informed throughout the process ● Monitors and reviews risk, risk treatments, obligations and criteria on an ongoing basis, identifying and responding appropriately to significant changes
ITIL	● Supports procurement and finance ● Scales quickly and reduces IT overcapacity ● Leverages new technologies ● Reduces IT ownership	● **Business relationship management**—Forms and upholds the cloud service provider and customer business relationship ● **Demand management**—Helps to understand, anticipate and influence business demand for services; calculates demand to allocate the agreed budget within the financial management process ● **Financial management for IT services**—Provides a cost-effective administration of the assets and resources used in providing IT services; incorporates charging based on consumption when perusing cost analysis calculation ● **Service portfolio management**—Describes provider services in terms of business value and needs; creates a portfolio for all potential external cloud deployment models ● **IT service management**—Assesses the service provider's offerings, capabilities and competitors, and current and potential market spaces to develop a strategy to serve customer needs

Source: Reijnders, B.; "A comparison of governance models for cloud computing," Tilburg University, http://arno.uvt.nl/show.cgi?fid=144876

Operational and Security Requirements

When building the cloud governance framework, there are two types of requirements that must be met: operational and security. Operational requirements focus more on aligning with business goals and setting and monitoring SLAs. Security requirements focus more on provider evaluation and the controls checklist.

Operational and security requirements may overlap at some point. This usually occurs when customers evaluate the requirements for cloud application availability, assessing risk based on the SLA, and on backup and disaster recovery (DR) plans.

Operational Requirements

Following are operational requirements:

- **Define goals for cloud migrations**—When a customer plans a move to the cloud, the first thing that needs to be done is to define the goals of the cloud migration. Different stakeholders have different goals from the migration.

These could be strategic goals (reducing costs, reducing environmental footprint, adding agility) or operational goals (increasing performance or availability).

- **Define the provider evaluation criteria (market share, reputation)**—The operational requirements include making sure that a provider is mature and well-established. This is accomplished by filtering providers based on their business and market status. Some organizations rely on analyst opinions (i.e., Gartner magic quadrant). Some simply consider size, income, number of employees, market share and other business indicators.
- **Define the internal SLA (system uptime, response time, downtime, backup retention)**—The SLA is the foundation for operational requirements and therefore requires careful review. An internal SLA can be helpful to determine the service level that the different units in the organization expect (e.g., the IT department can roll out a database server within four hours of a request) and compare them with what a CSP offers.
- **Define specific requirements (backup, monitoring, disaster-recovery procedures)**—In some cases, the provider does not cover all the operational requirements. For example, a CSP may offer internal backup services only, and the organization decision may be that all cloud services will be deployed on a public cloud. It is up to the governance process to set those specific requirements and monitor their successful implementation.

Security Requirements

Following are security requirements:

- **Define the provider evaluation criteria (provider certification, specific jurisdiction)**—Cloud customers reduce the risk of immature providers by excluding those that do not meet certain criteria, such as maturity level. In security, this filtering process usually happens by qualifying only providers that can demonstrate certain security baselines or gain certain attestation or certification.
- **Define the controls framework (using CCM, ISO27001 or similar)**—Organizations use a controls framework to verify that all relevant controls are in place and used properly. For each cloud migration or cloud application, the cloud customer should define the applicable controls and who is responsible for setting them.
- **Define specific requirements (location of encryption key, security information event management (SIEM) integration)**—Control frameworks may not have all the details, or cloud customers might require different specifications. As part of the security requirements process, the cloud customer defines specific requirements on top of the controls framework. For example, a personal data encryption key must always be kept in a hardware security module (HSM).

When building operational and security requirements, organizations also need to account for the legal and regulatory requirements that might affect those requirements.

Tip: The EU regulation for digital service providers and telcos provides a mechanism for incident monitoring and fines for service disruptions. Adhering to this regulation is highly dependent on the availability SLA providers offer.

Legal and Regulatory Requirements

Just like organizations set their security and operational requirements, they also need to set their legal and regulatory requirements. In some cases, the legal requirements may influence the operational requirements.

Tip: Incident reporting is an example of a legal obligation that may influence the SLA requirements. In some cases, the only way for the cloud customer to know about an incident is because the provider must report it. Some examples of laws and regulations with incident reporting:

- EU GDPR has a breach notification obligation.
- The EU digital service provider regulation requires incident reporting on service disruptions.

The process of setting legal and regulatory requirements involves the following steps:

1. **Understand the law and relevant regulations that apply to the organization**—This is done once for the entire organization, and it also needs to be done for each cloud migration or cloud workload.
2. **Classify data or operations that might require special attention**—For example, personal data may be kept in a specific jurisdiction, but applications with high availability requirements need to be deployed in two locations.
3. **Establish provider contractual negotiation guidelines**—Not all cloud providers will negotiate the agreement, and not all the clauses in the agreement are open for discussion. However, negotiation still needs to occur at some level (e.g., most organizations negotiate the location of dispute resolution).
4. **Set provider evaluation criteria**—Certain regulations may affect the provider selection process (e.g., the provider must have a specific geographic presence or must attain a particular certification).
5. **Understand requirements coming from contractual obligations**—For example, a customer contract may state that data will not move to a third party without consent.

The result of this process can be pure legal requirements (e.g., place of dispute resolution) or operational requirements (e.g., the provider must have two data centers at two remote geographic locations for availability). It is up to the cloud customer team to fuse those requirements into one large requirements matrix. See section 2.6 for more information on legal and regulatory requirements.

1.4.10 Cloud Risk Management

This section discusses cloud risk management.

Why Cloud Risk Management?

The rapid growth of cloud adoption and steady increase in global deployments significantly influence traditional technology risk management practices and the role of enterprise risk management (ERM) frameworks and programs. Some of these influences are a function of the CSP service and deployment models. Cloud providers deliver business-related and technology infrastructure services, and the risk associated with each requires identification and analysis throughout the risk management life cycle. In addition, the number of services continues to grow exponentially, vastly expanding the breadth and scope of coverage required by an ever-increasing number of CSPs. This rapid growth creates a risk due to the pure volume of third-party service providers that require various levels of initial and ongoing monitoring.

Examples of Cloud Risk

This section provides two examples of cloud risk.

Integration with ERM

A frequent organizational rationale for adopting cloud services is to rapidly implement new services and products designed to achieve core business or organization strategies and objectives. This alters the risk equation to include management and measurement of the opportunity risk associated with cloud success. This requires an organization dimension of the risk management function in the absence of an ERM.

In many organizations, cloud risk is directly related to strategic risk, conferring heightened visibility and importance to the entire organization—not just the business or functional owner of the CSP relationship.

New Inherent Risk Factors

"The cloud—with its complexity—is also the perfect place for attackers to hide. It is also, unfortunately, an ideal launchpad for attacks," according to a Cloud Security Alliance study.[10]

Although there may be an initial organizational tendency to view cloud governance, policies and evaluation methods as an extension of historical third-party practices, there are significant differences. Some major considerations include (**figure 1.13**):

- **Isolation failure**—Cloud computing often uses shared resources and multitenancy. There are security mechanisms to separate storage, memory, routing, etc., between different tenants. Failure of isolation mechanisms is considered a risk.

Example: Due to a hypervisor bug, it is possible to access the RAM of a neighboring cloud computer that belongs to another cloud customer. These types of software vulnerabilities are always a risk that can lead to unpredictable events and outcomes.

- **Interface compromise**—Public cloud providers enable their customer management interfaces to be accessed over the public Internet. If those interfaces are compromised by unauthorized external parties, the cloud customers' computing and data can also be compromised easily, through capabilities such as remote access and change of cloud security policies.

Example: A developer error of hardcoding cloud access credentials into a configuration file and an admin error of creating a public image of an application server can give hackers access to a tenant cloud management interface. These additional software glitches or threats—unique to the use of cloud technologies—need to be remediated.

- **Incomplete data deletion**—When a cloud customer makes a request to delete a cloud resource or data, it might not result in true and complete wiping of the data, because CSPs may have extra copies of data/resources for availability and reliability purposes.

Example: Due to the high availability goals of a CSP, the command to delete sensitive data may not result in full destruction. Also, due to the customer's limited visibility of the environment, the customer's command to delete the data might be inadequate. These are relevant compliance risk scenarios to consider, together with appropriate preventive or detective controls that the CSP adopts.

- **Shadow IT**—Cloud services from an online provider can be acquired and contracted as easily as ordering a book. This ease of acquisition spawned the growth of unmanaged and even unknown CSPs that provide technology resources, computer processing, and data communication and storage, while avoiding internal assessment and due diligence regimens.[11]

Example: It is very easy for a non-IT operational department head to subscribe to a SaaS solution in the cloud and start operations without the prior knowledge of the IT department. Once the risk is identified, the customer needs to establish acceptable preventive or detective controls.

- **Cloud supply chain**—Another significant difference between noncloud technology resources and platforms vs. cloud is the importance of other cloud platforms or service providers that may be transparent to the customer. Hence, it is up to the customer to use its discovery and risk assessment capabilities to identify the entire supply chain supporting any single cloud service provider. The more complex and extensive the supply chain, the

[10] Cloud Security Alliance, "Top Threats to Cloud Computing: Egregious Eleven," 6 August 2019, https://cloudsecurityalliance.org/artifacts/top-threats-to-cloud-computing-egregious-eleven

[11] Dimicco, R.; "Shadow IT: Rampant, Pervasive, and Explosive!", Cisco Blogs, 19 January 2016, https://blogs.cisco.com/cloud/shadow-it-rampant-pervasive-and-explosive

greater the inherent risk that one or more of the individual providers can jeopardize the customer's compliance with relevant regulatory and legal requirements.

- **Cloud configuration risk/opportunities**—The very nature of cloud computing—the variation of service and deployment models, the ownership and accountability for all controls without responsibility for all of them—presents new and challenging tasks for cloud customers. The abstract nature of operating in a third-party data center requires means of gaining visibility into data and information that are not straightforward or may not be readily available. Much of this information involves data and metadata regarding configuration settings and their management.

Figure 1.13—New Inherent Cloud Risk Factors

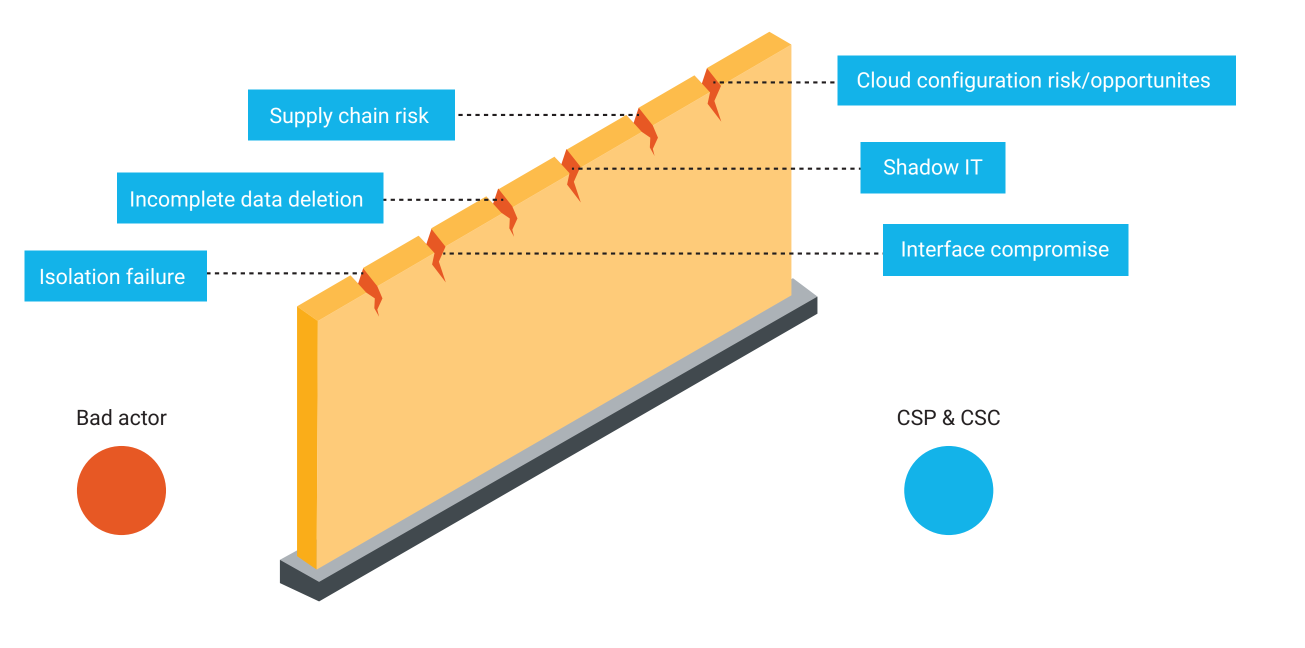

Understanding Risk Management

Although cloud computing mitigates some risk, it also creates new technical, operational and compliance risk. Therefore, it is important for cloud customers and CSPs to perform continuous risk management, which requires an understanding of the details and different approaches to risk management implementation.

Risk Management (Definition)

The International Organization for Standardization (ISO) defines risk as "effect of uncertainty on objectives."[12] It defines uncertainty as a "state, even partial, of deficiency of information related to understanding or knowledge of an event, its consequence, or likelihood."[13]

Risk is the potential for unforeseen loss of something of value. In cybersecurity, the value is often referred to as the asset, which can be the data or the computing capability. Loss of value can initially be confidentiality, integrity or availability of any data or capability. Later, that loss can damage the organization objectives. In cloud environments, the term value can be the cloud environment itself, the data that are processed and stored in the cloud, and the effect of the data management on the stakeholders and business objectives.

[12] ISO, ISO GUIDE 73:2009, *Risk management – Vocabulary*, 2016, www.iso.org/standard/44651.html
[13] *Ibid.*

The ISO guide defines risk management as "coordinated activities to direct and control an organization with regard to risk."[14] Like any other emerging technology, cloud computing creates new risk that cloud customers and CSPs should address.

Risk Management Approach

Risk management includes many processes. The European Union Agency for Cybersecurity (ENISA) groups those processes into three main domains and two supporting domains as shown in **figure 1.14**.

Figure 1.14—ENISA Risk Management Process

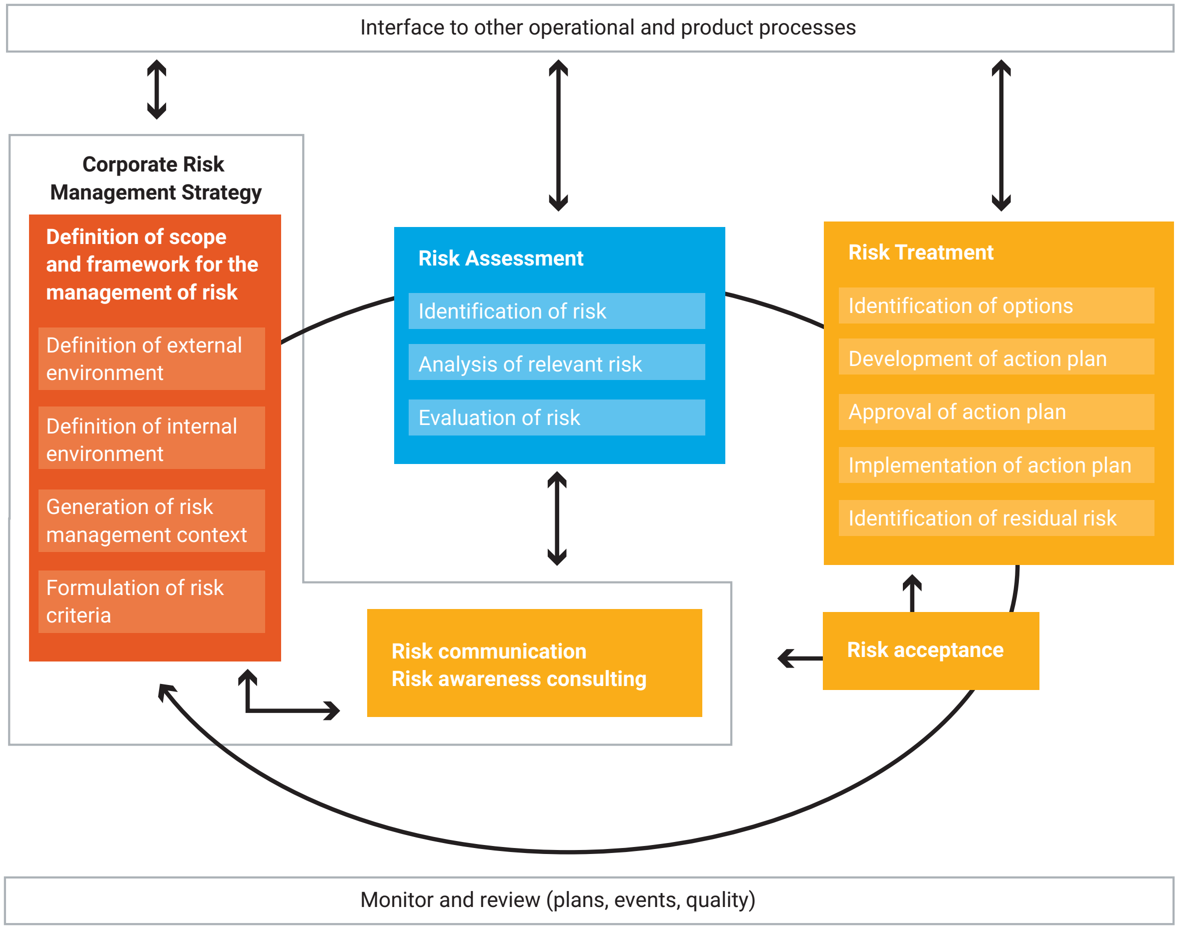

Interface to other operational and product processes
Corporate Risk Management Strategy
Definition of scope and framework for the management of risk
Definition of external environment
Definition of internal environment
Generation of risk management context
Formulation of risk criteria
Risk Assessment
Identification of risk
Analysis of relevant risk
Evaluation of risk
Risk Treatment
Identification of options
Development of action plan
Approval of action plan
Implementation of action plan
Identification of residual risk
Risk communication
Risk awareness consulting
Risk acceptance
Monitor and review (plans, events, quality)

Source: ENISA, "Risk Management Process," www.enisa.europa.eu/topics/threat-risk-management/risk-management/current-risk/risk-management-inventory/rm-process/rm-process

Strategy/Scope and Framework

ENISA recommends four important steps to create a sound strategy, scope and framework for risk management:

[14] *Ibid.*

- **Definition of internal environment**—The details of the internal environment, internal stakeholders, existing risk, organizational structure and culture, important assets, goals and objectives, and existing strategies help build an understanding the internal environment, which is used to select and customize the right risk management approach.

Example: Organizations in the healthcare industry operate in environments with the inherent risk associated with safeguarding protected healthcare data (PHI). Their environments need appropriate mechanisms to identify, classify and protect all such data to ensure regulatory compliance.

- **Definition of external environment**—The external environment typically includes the market of reference, the competition and other external stakeholders, including political, socio-economic, financial, legal and regulatory environments.

Example: Organizations working in a country with frequent changes of government might be subject to changing of relevant laws and regulation, which might affect (both directly and indirectly) the business (e.g., new antitrust laws, new privacy laws, etc.).

- **Generating the risk management context**—The business case of risk management provides a balance between costs, benefits and opportunities.

Example: Enterprise risk management is a cost of doing business when adopting cloud computing. It must apply to all applications, systems, data being processed, stored and communicated, and to all personnel who have access to such systems and data stores. The resulting risk management program should be designed to meet and not exceed the risk appetite of the organization.

- **Formulation of risk criteria**—As the risk register (the list of all the risks and their relative importance) grows, it is important to have a fixed risk formulation and criteria to calculate each risk's priority. During this step, risk should be assessed using criteria, such as likelihood—i.e., chance (usually a percentage that the risk will become reality) or impact (the measure, often in currency, of the possible effect of the risk becoming a reality).

Example: Regardless of the organization size, it should consider each risk for probability of occurrence and potential impact. The totality of all mathematical results should not exceed the organization risk appetite. The challenge in most organization scenarios is quantifying both probability and impact.

- **Communication, awareness and consulting**—Risk management can be successful only if it is embedded in the organization and is clearly communicated to all stakeholders.

Example: Communication and operation of the risk management program must be built into the fabric of the organization to maximize effectiveness. For adopting new and emerging technologies, such as cloud computing, extra efforts are required to create pervasive awareness and training programs throughout the organization.

Risk Assessment

Risk assessment is the process used to identify and evaluate risk and its potential effect within the risk management domain, which includes subprocesses relevant to discovering and determining, understanding, prioritizing, and documenting the risk. Every cloud offering and use case is subject to risk, and risk assessment focuses on identifying, analyzing and evaluating risk to treat them appropriately.

Identification of Risk

The first step in risk assessment is to identify the specific risk. Risk identification is the process of determining risk conditions that could potentially delay implementation of the cloud solution or prevent it from achieving its objectives.

The two main questions to ask when identifying risk: What can happen? And why can it happen? This step can also include identification of the risk owner. In cloud environments, the risk owner can be any group within the cloud customer, CSP, or other party in the cloud supply chain.

Example: An organization is using a cloud-based SaaS customer relationship management (CRM) software application. Besides generic CRM data, the organization is also storing top-secret customer information within that solution. The identified risk can be that top-secret data can be disclosed to unauthorized parties. The risk owner is the owner of the top-secret data. This is the inherent risk, so the organization needs to consider the controls that are in place, e.g., encryption.

Analysis of Relevant Risk

Risk analysis is the phase in which level of risk and its nature are assessed and understood.

Cybersecurity-related risk consists of three components: actor, threat and vulnerability.

- A threat actor is an entity (individual, group, organizations or state) that seeks to exploit the organization dependence on cyberresources.
- Threat is a possible danger that can harm the data or capability. Threats can be human or digital in nature.
- Vulnerability is a weakness in a process or system that can be exploited by a threat. Vulnerabilities are usually internal or inherited by the target and caused by a missing component or capability.

After defining the main components (actor, threat and vulnerability), the next step is defining the potential for loss, damage or destruction of an asset due to a threat exploiting a vulnerability.

Analysis is the process of breaking down the risk into its components and identifying applicable controls (e.g., security precautions).

See chapter 4 for more information on factors, threats and vulnerabilities.

Example: In the previous scenario, there is a risk that the top-secret data in the CRM tool can be exposed to unauthorized parties, internal or external. This possibility requires the risk assessment process to appropriately reduce the risk.

Evaluation of Relevant Risk

Organizations need to prioritize risk, so they can use limited resources to treat them optimally. The risk priority (i.e., score) is often calculated as a function of its impact (potential damage) and its likelihood (probability of becoming a reality). Based on the risk assessment methodology that an organization uses, the processes may include some other attributes of the risk.

Example: Continuing with the previous scenario, the impact of unauthorized disclosure of top-secret data in SaaS CRM applications is high, and its likelihood of occurrence is also high, so the organization records this risk in the risk register with the relevant priority. This drives the amount of remediation attention it receives and required communication to stakeholders.

Risk Treatment

After risk conditions are assessed, it is time to treat them. There are four different methods for the risk treatment domain: risk avoidance, risk reduction or mitigation, risk transfer and risk acceptance. The risk treatment process includes:

- **Identification of options**—Based on the risk components (i.e., threats and vulnerabilities) and the asset in question, find relevant options to treat the risk.

- **Development of action plan**—Based on the identified options, select (and group, if required) different options into an action plan.
- **Approval of action plan**—The risk owners review and approve the action plan.
- **Implementation of action plan**—Implementation of the action plan may be instantaneous or require a long-term project to perform changes.
- **Identification of residual risk—**Identify risk after the action plan is implemented.

Risk Avoidance

Risk avoidance is the elimination of threats, activities and vulnerabilities that can cause a risk.

Example: The risk owner can decide to move the top-secret data to a more secure application to avoid the risk. This also requires validation that the more secure application meets expected control capabilities in both design and operation.

Risk Reduction

Risk reduction is the process of finding ways to take on less risk, including avoiding execution of related activities.

Example: Continue using the CRM but do not process top-secret data in that SaaS solution. This option requires confirmation that the data that remains in the CRM do not pose a risk that exceeds the organization risk appetite.

Risk Mitigation

Risk mitigation is reducing the likelihood or the consequences (damage) of the risk. This is done by taking countermeasures to reduce the level of the risk.

Example: Improve the SaaS solution to include two-factor authentication by requiring that MFA codes be sent to a trusted mobile device after successful username and password entry. This additional control has to be validated for its ability to lower the total data risk to within the organization risk appetite.

Risk Transfer

Risk transfer is a strategy of dealing with risk by transferring it to another person or entity, such as an insurance agency. With proper contracts, moving to the cloud can also be a risk transfer method. The shared responsibility model of the cloud transfers some risk to the CSP. Risk transfer techniques need careful consideration before being accepted as an option. For example, using insurance as a risk transfer requires an assessment by the insurance company as a basis for setting the premiums. Any losses require filing a claim, which the insurer will not pay unless the insured complied with all the terms and conditions of the policy. Moreover, the benefit is only financial. The organization can still be liable for regulatory fines and penalties and derivative lawsuits, and suffer the reputational risk associated with the event.

Example: Risk owners can get an insurance policy against cybersecurity risk. It should be recognized that the purchasing of insurance is not a panacea for risk reduction. Cybercoverage typically requires a review and examination of risk and controls by the insurance company. If it provides coverage, the premiums it charges reflect the perceived quality of the control systems. In the event of a claim, any type of customer failure or negligence may obviate the claim. Moreover, even if the claim is paid, the organization still bears the brunt of reputational risk, exposure to regulatory fines and penalties, and shareholder or other stakeholder lawsuits.

Risk Acceptance

Unless the risk is avoided completely and the risk cannot be eliminated after all other selected risk treatment methods are implemented, there is residual risk. The risk owner has two options to deal with this: Either select one or more additional risk treatment methods (usually when the residual risk level is higher than the risk acceptance criteria), or accept the residual risk (applicable when the residual risk level is lower than the risk acceptance criteria). Once accepted, the risk owner agrees to live with the risk and face the residual consequences if the risk becomes a reality, despite the reduced likelihood. A risk acceptance decision is often triggered when the resources required to establish and operate a suitable control to mitigate the remaining risk are greater than the potential damage caused by the risk itself. Any mature enterprise risk program requires the accumulation and quantification of all accepted risk. This is an important number to manage and understand, to estimate the total exposure the organization may face based upon the nature of an organization-wide loss.

Example: Having an insurance policy against cybersecurity risk may not cover all the remaining risk against reputation loss after an incident. Yet management may decide to accept this residual risk. There are other factors to consider: 1) Typically, a single business executive has to take ownership of the risk acceptance, but once monetized, the executive may not be willing to do so; and 2) in mature risk management programs, all accepted risk is combined and totaled for a consolidated view of impacts. If more than one accepted risk condition can occur at the same time, acceptance may not be an appropriate decision.

Monitor/Metrics and Review

Risk management is a cyclically executed process that involves a set of coordinated activities for overseeing and controlling risk. Monitor and review include monitoring the risk and security controls in the cloud solution, assessing their effectiveness, documenting changes, and reporting risk levels and security states to the stakeholders.

Example: Due to the unique nature of the risk emanating from cloud computing and the rates of change introduced by the supporting technologies and infrastructure, the monitoring of key risk, controls and processes may warrant a need to be continuous. This increases costs, reduces discovery and remediation time, and requires the creation of relevant metrics and metadata. This becomes one of the most critical components of an effective ERM.

Cloud Risk Management Best Practices

This section presents best practices for risk management.

Risk Frameworks

It is possible to adapt and use legacy frameworks for cloud services. Following are some well-known standards and their scope and approach to risk management:

- **NIST 800-30**—The US Department of Commerce, National Institute of Standards and Technology (NIST) published its *Guide for Conducting Risk Assessments* (NIST 800-30) initially in 2002. The guidelines include a detailed taxonomy on risk and a definition of the risk management process.[15]
- **ISO/IEC 27005:2018**—Developed by the International Organization for Standardization (ISO) and the International Electrotechnical Commission (IEC), the document supports the general concepts specified in

[15] National Institute of Standards and Technology, "Guide for Conducting Risk Assessments," September 2012, https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-30r1.pdf

ISO/IEC 27001 and is designed to assist the satisfactory implementation of information security based on a risk management approach.[16]

- **FAIR**—The Factor Analysis of Information Risk (FAIR™) quantitative risk analysis model defines the necessary building blocks for implementing effective cyberrisk management programs.[17]
- **OCTAVE**—Developed by Carnegie Mellon University Software Engineering Institute, The Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE) is a framework for identifying and managing information security risk.[18]
- **EBIOS**—EBIOS Risk Manager (EBIOS RM) is the method for assessing and treating digital risk published by the National Cybersecurity Agency of France (ANSSI) with the support of Club EBIOS. It provides a toolbox that can be adapted, its use varying according to the objective of the project. EBIOS Risk Manager is compatible with the reference standards in effect, in terms of both risk management and cybersecurity.[19]
- **OCTAGON**—The Cloud Octagon Model by CSA makes it easier for organizations to identify, represent and assess risk in the context of their cloud implementation across multiple actors (legal, information risk management, operational risk management, compliance, architecture, procurement, privacy office, development teams and security).[20]

Additional Frameworks/Standards to Guide Cloud Risk Management

Although they are not risk management frameworks, the following documents provide invaluable risk and control guidance:

- **Cloud Controls Matrix (CCM)**—The Cloud Security Alliance created CCM to provide fundamental security principles to guide cloud vendors and to assist prospective cloud customers in assessing the overall security risk of a cloud provider. While CCM provides a controls framework in 13 domains, it also cross-references the controls to other industry-accepted security standards, regulations and controls frameworks, such as ISO 27001/27002, ISACA COBIT, PCI, NIST, Jericho Forum and NERC CIP.
- **Consensus Assessments Initiative Questionnaire (CAIQ)**—CAIQ is based on CCM. It provides a set of yes/no questions a cloud customer and cloud auditor may wish to ask of a CSP to ascertain its compliance with the Cloud Controls Matrix. Although CAIQ is not a suitable tool for risk analysis, it provides a clear and simple approach for assessment and auditing purposes.
- **Top Threats Deep Dive**—The Cloud Security Alliance created the Top Threats Deep Dive case study to consolidate and better understand all facets of security analysis by using nine anecdotes cited in the Top Threats for its foundation. Each of the nine examples are presented in the form of a reference chart and a detailed narrative.[21]
- **NIST CSF**—The NIST *Framework for Improving Critical Infrastructure Cybersecurity* (NIST Cybersecurity Framework, or CSF) was originally published in 2014.[22]

[16] ISO, ISO/IEC 27005:2018 *Information technology – Security techniques – Information security risk management*, July 2018, www.iso.org/standard/75281.html

[17] FAIR Institute, "FAIR Risk Management," www.fairinstitute.org/fair-risk-management

[18] Alberts, C.; S. Behrens; R. Pethia; W. Wilson; "Operationally Critical Threat, Asset, and Vulnerability Evaluation (OCTAVE) Framework, Version 1.0," Carnegie Mellon University Software Engineering Institute, September 1999, https://resources.sei.cmu.edu/library/asset-view.cfm?assetid=13473

[19] ANSSI, "EBIOS Risk Manager – The Method," November 2019, www.ssi.gouv.fr/uploads/2019/11/anssi-guide-ebios_risk_manager-en-v1.0.pdf

[20] Cloud Security Alliance, "Cloud Octagon Model," 24 June 2019, https://cloudsecurityalliance.org/artifacts/cloud-octagon-model

[21] Cloud Security Alliance, "Top Threats to Cloud Computing: Deep Dive," 8 August 2018, https://cloudsecurityalliance.org/artifacts/top-threats-to-cloud-computing-deep-dive/

[22] National Institute of Standards and Technology, "Cybersecurity Framework," www.nist.gov/cyberframework

Assessing the Cloud Supply Chain Risk

If the CSP outsources parts of its infrastructure, operations or maintenance, these third parties may not satisfy or support the requirements that the CSP is contracted to provide to cloud customers. An organization needs to evaluate how the CSP enforces compliance and check if the CSP flows its own requirements down to third parties. Having regular discussions with the CSPs on supply chain contractual requirements and activities through risk/KPI reports helps to identify risks that need mitigation. If the requirements are not being levied on the supply chain, then the threat to the customer increases. This threat increases as an organization uses more CSP services, and it is dependent on individual CSPs and their supply chain policies.

Shared Responsibility Model and Cloud Risk

Apart from some customers self-hosting in the cloud, cloud use includes some form of supply chain relationship. Depending on the cloud model used (e.g., IaaS or SaaS), the ownership and operation of the layers change. However, overall risk still exists, and the cloud customer should manage it. **Figure 1.15** shows the shared responsibility model and cloud risk.

Figure 1.15—Cloud Customer and Provider Shared Responsibility Risk Matrix

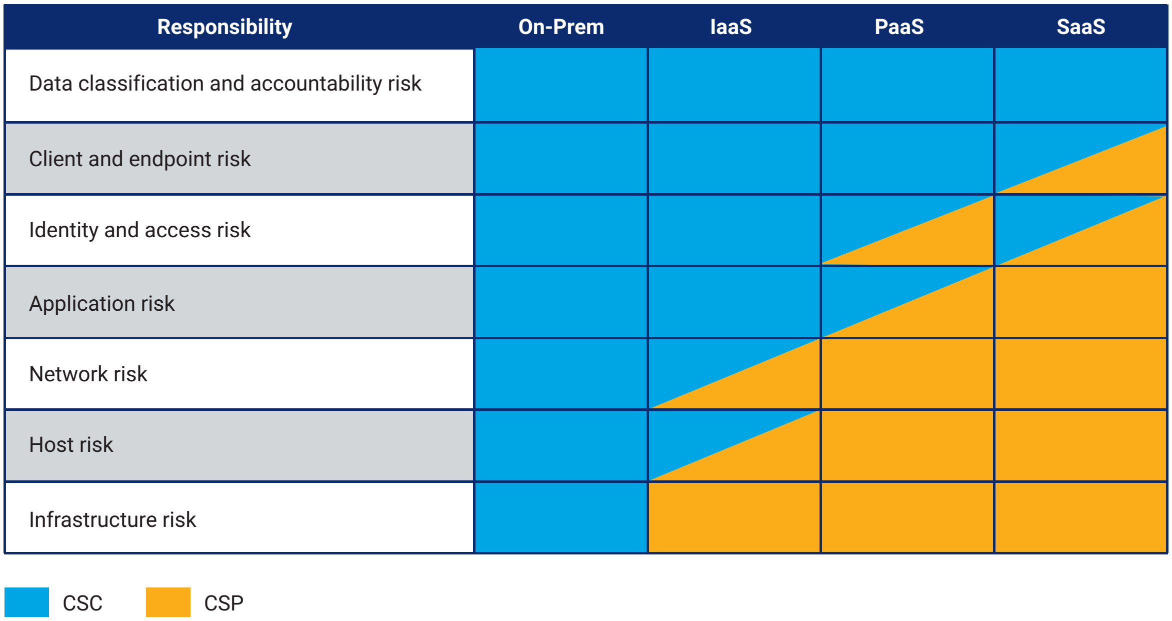

Responsibility	On-Prem	IaaS	PaaS	SaaS
Data classification and accountability risk	🟦	🟦	🟦	🟦
Client and endpoint risk	🟦	🟦	🟦	🟦🟧
Identity and access risk	🟦	🟦	🟦🟧	🟦🟧
Application risk	🟦	🟦	🟦🟧	🟧
Network risk	🟦	🟦🟧	🟧	🟧
Host risk	🟦	🟦🟧	🟧	🟧
Infrastructure risk	🟦	🟧	🟧	🟧

🟦 CSC 🟧 CSP

Source: Microsoft TechNet, "Shared Responsibilities for Cloud Computing," 25 October 2019, https://gallery.technet.microsoft.com/Shared-Responsibilities-81d0ff91

The CSP owns and is accountable for some risk in the cloud. Being the data owner, the cloud customer is ultimately accountable for all the risk throughout the stack, regardless of the cloud model it uses. Even if third parties are responsible for some tasks, cloud customers are still accountable.

1.4.11 Cloud Compliance

This section introduces cloud compliance and provides an overview.

Introduction to Compliance

In general, compliance means conforming to an order, rule or request, such as a specification, policy, standard, contract or law. In the context of cloud security, compliance refers to conforming to security-related requirements set by an organization policy, standard, controls or regulatory requirements to protect business objectives and infrastructures.

Examples:

- Performance specifications, which define the minimum number of transactions performed per minute
- Security policy, which requires two-factor authentication to access customer data
- PCI DSS standard, which requires that credit card information be stored with encryption
- Customer contract, which orders wiping of the data after processing
- Privacy law, which prohibits sensitive data from leaving a jurisdiction area

In legacy IT, the organization has the sole ownership of those compliance requirements. The accountability, identification, planning, execution, reporting and assurance of those requirements belong to the organization, and it cannot delegate them. However, the organization can transfer risk to a responsible third party, consulting and informing others as required.

How Cloud Impacts and Changes Compliance

The International Organization for Standardization (ISO) compliance management system guidance document[23] defines compliance as an outcome of an organization meeting its obligations. Compliance is made sustainable by embedding it in the culture of the organization and in the behavior and attitude of people working for it. Compliance is usually measured once per year in the context of a SOC 2 type II audit. Every day, the IaaS/PaaS compliance dashboards present compliance data to the team (e.g., risk or compliance team, or management). In an audit, compliance of controls are tested over a longer period.

Compliance in the cloud should be considered from the perspectives of the CSP and the cloud customer, as follows:

- **Customer**—The customer is ultimately responsible for all its compliance requirements. In addition, the customer must be sure that any cloud usage does not jeopardize the organization's compliance efforts.
- **CSP**—The CSP is responsible for its own compliance requirements. A CSP also supports the customer by providing compliance details and any compliance certifications for the related service. This approach (i.e., compliance inheritance) helps the customer to include cloud services in its own compliance efforts.
- **CSP as a customer**—A CSP may not be the sole provider of a cloud service. It can be part of multilevel supply chain structure that includes other (sub) cloud services. The CSP can be a provider and a customer at the same time, which can result in multilevel compliance inheritance.

Example: A well-known SaaS video-on-demand provider, Netflix, uses another PaaS and IaaS cloud provider, Amazon Web Services (AWS).[24] Any relevant Netflix compliance requirements should be communicated to AWS.

[23] ISO, ISO 19600:2014(en) *Compliance management systems – Guidelines*, www.iso.org/obp/ui/#iso:std:iso:19600:ed-1:v1:en
[24] AWS, "Netflix on AWS," https://aws.amazon.com/solutions/case-studies/netflix/

Cloud computing has a considerable impact on the compliance posture of CSPs and customers. From the CSP perspective, the main challenge is possibly the complexity of the legal and regulatory framework (e.g., privacy rules vary by region and countries). Given the global nature of most CSPs' offerings and footprints, sometimes they must reconcile compliance with conflicting laws and regulations.

Example: The EU General Data Protection Regulation (GDPR) and USA CloudACT conflict with each other. While the CloudACT specifically contemplates court orders and warrants requiring the transfer of personal data without an international agreement, the European regulators concluded that "service providers subject to EU law cannot legally base the disclosure and transfer of personal data to the US on such requests."[25]

In addition, CSPs are often expected to comply with multiple national, international and sector-specific laws or standards, which puts them under considerable pressure and exposes them to the issue of compliance complexity and fatigue (see section 2.10 for more information). From the customer's perspective, there are similar challenges related to the complexity of the legal and regulatory landscape in addition to challenges related to the shared responsibility model. Like with cloud security responsibilities, a cloud customer can offset some compliance responsibilities to the CSPs (see "Compliance Inheritance" section), but it cannot outsource compliance accountability.

Compliance continuity is the concept of maintaining compliance without interruption throughout any planned or unplanned change or event, whether internal or external. When moving to the cloud, the customer must take steps to ensure continuity of compliance. Planning for continuity of compliance with any cloud-related change throughout the cloud journey is key. This may include development of cloud specifications, submission of requests for information (RFIs) or requests for proposals (RFPs), CSP selection, service transition, service operations, and continuous monitoring of CSPs' compliance efforts and changes.

Evidence is crucial to any compliance audit or internal controls process because it affords signatories a reason to trust outstanding claims. Compliance management includes collecting evidence to show how the CSP supports the specific compliance rules and requirements.

The cloud customer should reiterate all the compliance rules, orders and requirements because some inherited compliance rules might be forgotten over time.

Example: A local business that has been using a local data center may have avoided data export restrictions in the past. When the business decides to move to a public cloud provider, the rules and regulations, and the CSP options should be iterated first.

The responsibility for cloud security and compliance is shared. The CSP and the cloud customer have a role to play.

Scope of Compliance and the Organizational Accountability Model

With the introduction of the cloud, customers transfer to CSPs not only some IT stacks, but also some responsibility for policies, standards, business requirements, and legal and regulatory requirements.

Although security is a shared responsibility between the cloud service provider and the organization, with responsibilities distributed across the stack, compliance accountability remains with the customer.

Generally, compliance responsibility reflects the degree of control that a party has over the architecture stack. **Figure 1.16** shows a general shared responsibility model depending on the architecture.

[25] European Commission, "Data Protection in the EU," https://ec.europa.eu/info/law/law-topic/data-protection/data-protection-eu_en

Figure 1.16—Cloud Shared Responsibility Model

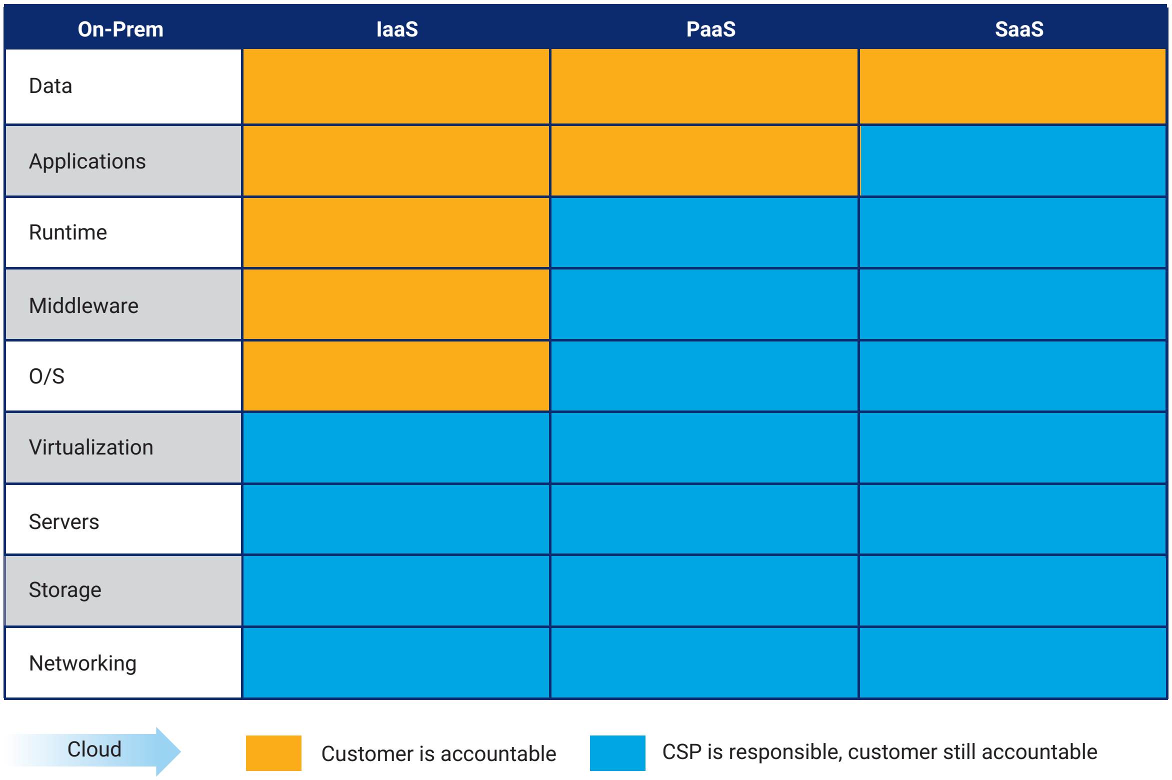

On-Prem	IaaS	PaaS	SaaS
Data	🟧	🟧	🟧
Applications	🟧	🟧	🟦
Runtime	🟧	🟦	🟦
Middleware	🟧	🟦	🟦
O/S	🟧	🟦	🟦
Virtualization	🟦	🟦	🟦
Servers	🟦	🟦	🟦
Storage	🟦	🟦	🟦
Networking	🟦	🟦	🟦

Cloud → 🟧 Customer is accountable 🟦 CSP is responsible, customer still accountable

Source: Scott, T.; "CSRM – Cloud Shared Resource Model," Open Alliance for Cloud Adoption, 13 October 2018, www.oaca-project.org/2018/10/13/csrm-cloud-shared-resource-model/

While the shared responsibility model is more straightforward for public and private clouds, a hybrid cloud inherently carries more fluidity, because it encompasses various deployment and service models.

Cloud customers and CSPs may rely on several other actors to carry out business activities. This also affects compliance responsibilities, as follows:

- **Cloud customer**—Many actors may be involved in the compliance efforts. Internal (e.g., different business units, information-communication technologies, information-cybersecurity systems) and external (e.g., managed service providers, business partners, auditors) parties can be part of the cloud customer.
- **Cloud service provider**—Many actors also can be part of CSP operations, such as subcontractors and subcloud service providers.

When the supply chain consists of multiple layers (e.g., a SaaS CSP uses a different IaaS CSP to deliver the service), each layer interacts only with the adjacent entity (at least legally). The cloud customer should ensure that all possible compliance issues inherent to the CSP subcontractors are under control.

Compliance Inheritance

In cloud computing, compliance inheritance refers to the idea of a CSC taking advantage of the CSP compliance efforts. Typically, cloud customers inherit compliance controls from a CSP in areas that are closely related to the security of the cloud infrastructure. Viewed more generally, the shared responsibility model affords insight into which controls are inherited from the CSP.

Having such a mechanism in place is extremely important, because most public cloud providers do not grant customers the right to first-party auditing for compliance assurance.

Example: The Guidelines on Securing Public Web Servers (NIST 800-44 Version 2:2007) suggest that web servers be located in areas that provide secure physical environments with appropriate physical security protection mechanisms in place, including appropriate environmental controls to maintain the necessary humidity and temperature, a backup power source, etc. When an organization using a cloud solution for its web servers is required to show its alignment with this NIST document, the audit scope should include the physical security controls.[26]

A common mechanism for compliance inheritance concerns relevant third-party audit reports, such as certifications, which allow cloud customers to leverage attestations to fulfill compliance due diligence. The provider may have parts of its service certified as compliant, which would remove it from the customer's audit scope. Compliance inheritance is effective when a CSP has the capabilities to demonstrate compliance—including document generation, evidence production and process compliance—in a timely manner.

Example: Continuing with the previous example, even if the cloud provider does not allow any first-party audit, an independent CSA STAR, ISO 27001, PCI DSS or similar certification including related physical security controls could be used to show compliance.

Cloud customers can be challenged to show auditors that the organization is in compliance. Assignment of compliance responsibilities to the provider and customer, and to indirect providers (i.e., the CSP of your CSP), can help both the cloud customer and the auditors.

The scope of compliance plays a crucial role (**figure 1.17**). With compliance inheritance, the CSP infrastructure is out of scope for a customer's compliance audit, but everything the customer configures and builds on top of the certified services is still within scope. Building on top of a compliant CSP does not necessarily mean that the cloud customer is compliant too.

Example: Continuing with the previous example, collection of personally identifiable information is a function of the web application databases that the cloud customer manages and administers. It is not possible for the IaaS CSP to provide compliance with the privacy requirements of NIST 800-44 on behalf of the web solution. Hence, use of CSA STAR certified IaaS cannot solely ensure the customer's compliance.

[26] Tracy, M.; W. Jansen; K. Scarfone; T. Winograd; "Guidelines on Securing Public Web Servers,“ National Institute of Standards and Technology, September 2007, https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-44ver2.pdf

Figure 1.17—Compliance Inheritance

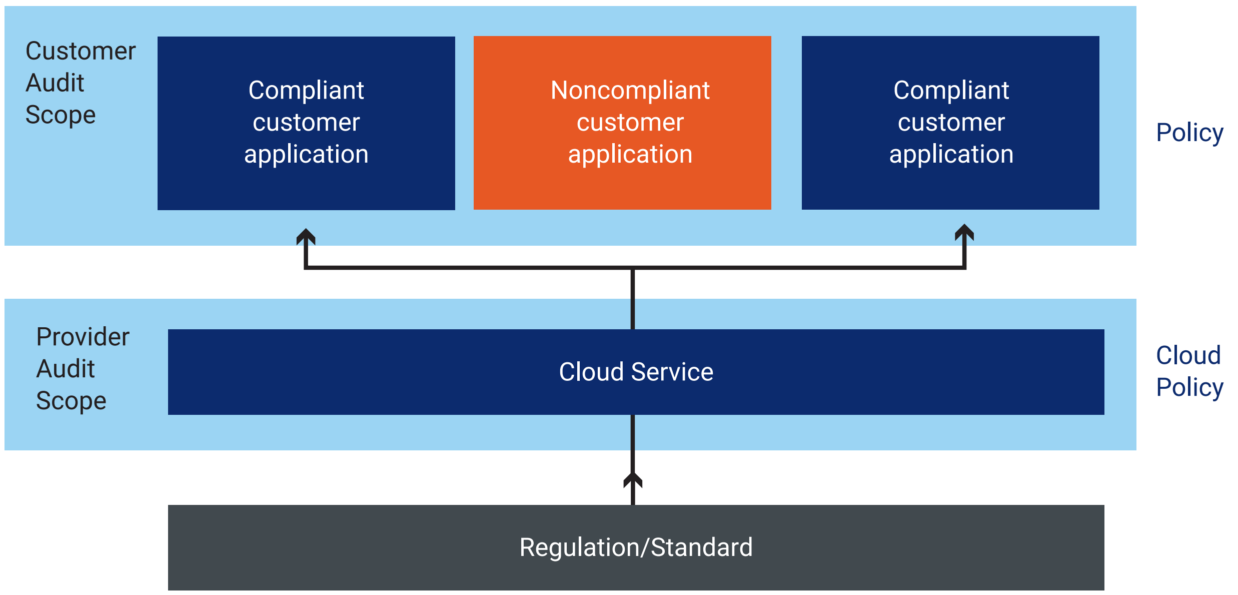
Customer Audit Scope
Compliant customer application
Noncompliant customer application
Compliant customer application
Policy
Provider Audit Scope
Cloud Service
Cloud Policy
Regulation/Standard

Source: Mogull, R.; J. Arlen; F. Gilbert; A. Lane; D. Mortman; G. Peterson; M. Rothman; "Security Guidance for Critical Areas of Focus in Cloud Computing v4.0," Cloud Security Alliance, 2017, https://downloads.cloudsecurityalliance.org/assets/research/security-guidance/security-guidance-v4-FINAL.pdf

When an organization uses any external service (such as a CSP), the external party is required to be compliant to the same orders, rules and requirements as long as it is part of the related data processing.

1.4.12 Cloud Governance Tools: Cloud Policy Enforcement

This section covers all aspects of cloud policy from development to enforcement—either for a distinct standalone cloud policy or an integrated one reflected across an existing set of policies. The latter is likely to be the more common case.

Cloud Policy Definition

Policies, standards and guidelines are related documents, although each serves a different purpose. For example, policies typically express at a high level how a particular risk is managed. According to the ISACA *CISM Review Manual*, "policy is a statement of management intent." Supporting standards usually define controls to achieve the goals of the policy. Guidelines describing how may be used to further supplement the standards. What one organization calls a policy another calls a standard. For example, a bank may use policies to express, at a high level, the what of how it manages a particular risk. It may have supporting standards to define controls to achieve the goals of the policy. It may further supplement the standards with guidelines that describe how. It is important for the customer to identify quickly how the organization uses those terms and whether it uses any recognized policy frameworks. If it is using a framework developed in-house, it is important to understand where and how it is referred to in the risk management framework when assessing it.

The intent of policies is to define the principles, criteria and approaches to business risk management and operational improvement applicable to the organization. At the highest level, policy is a description of a set of control objectives or targets established to ensure that policy intent is met.

Cloud policy statements often are woven into existing organization policy frameworks. Cloud computing is a relatively new concept, and most organizations are on a journey to get there. The rules an organization applies to its cloud operations may be spread out among its existing policies, standards and guidelines rather than spelled out in a specific named cloud policy. For example, rules on access control and authentication for accessing confidential information will most likely be in the broader policy framework.

First, the auditor should look at the cloud policies, and in their absence make the following determinations:

- The organization has identified cloud security as a risk.
- It has evaluated how to manage that risk.
- It has at least a work-in-progress set of control objectives that respond to that risk.

A useful first step in understanding cloud security policy objectives or statements is to have a discussion with someone in the organization who is familiar with its policy framework and can discuss how any cloud-specific risk or control objectives have been addressed. If those elements are clearly absent, the auditor can call that out. However, if there are statements applicable to the cloud in other policies, it is not necessary to recommend that the organization has a standalone or separate cloud policy.

Where Does Policy Fit?

Organizations typically have a hierarchy of policies to support their enterprise risk management. Some may have an overarching enterprise risk management framework (ERMF) document that describes how they identify, evaluate, manage and respond to various risk conditions encountered in the process of doing business. Other organizations may have higher level governance policies, with the ERMF considered a subcomponent.

The ERM will likely have a series of policy documents that speak to particular risk named in the framework—such as a technology policy, which might fall under operational risk. Typically, cloud policy-related statements fall under those headings. Each policy statement has a goal and often refers to standards that define actions that support the goal. Under the standards, there may be optional guidelines that drill down to a procedural level—e.g., how the organization can deliver on what is defined in the standard.

Breakdown of Policy Life Cycle

With CCM, the CSA is standardizing the language applicable to policy-related control objectives. Specifically, the risk management policy life cycle involves establishing, documenting, approving, communicating, implementing, enforcing and maintaining policies and procedures.

- **Establish**—Having determined that it needs a policy, the organization defines its scope and goals, and considers how it might align with other policies.
- **Document**—After it has determined what the controls should be, the organization goes through a drafting process to develop a documented policy for approval. It also considers whether the policy needs to meet any regulatory or legal requirements.

- **Approve**—Organizations often have a policy approval board or similar entity charged with reviewing proposed policies for the following:
  - Fit
  - Potential impact
    - Does it materially increase the cost of doing business?
    - Does it impact other functions?
    - Are there provisions for exceptions to the policy?
  - Alignment with other policies
  - Any missing material items
  - Adoption issues
    - Does it require further clarification?
    - Does it need one or more supporting standards to underpin it?
    - Does it have an implementation period?

  The approval board also checks evidence of consultation with various stakeholders.
- **Communicate**—The organization decides how to make it known that a policy has been approved (this is often a process, not an event). The information can be relayed as a communication to all employees, highlighted to senior managers to encourage the right behaviors regarding the policy. In a regulated environment, it may be communicated to persons with formal regulatory responsibility. There may be reminders, compliance dates, or deadlines to drive awareness and compliance with the policy, and to ensure that any concerns can be raised promptly.
- **Implement**—During the implementation period, program managers are assigned to develop an implementation plan or a program of work that is typically project-managed. During this phase, there may be calls for supporting standards, clarifications, additional budget allocations or personnel. There may be exceptions raised (although this may not begin until the enforcement phase). The policy creators may be contacted frequently during this phase. This is when the policy in its ideal state undergoes real-world implementation.
- **Enforce**—The enforcement phase typically follows the implementation period. Designated control owners are expected to have implemented the policy and be operating the controls in accordance with the policy. The flow of exceptions is lower than during the implementation period. Once a policy is being enforced, it is subject to planned or random internal audits. There may be measures or metrics showing evidence of compliance, and the policy itself may place evidential requirements on the control owners or operators.
- **Maintain**—The maintenance period is a period of reflection. As the organization and regulations evolve and as threats change, there is a need to update policies. An organization typically has a predetermined schedule for reviewing policies. The policy owner may state there are no changes required, or that a minor update is needed, or that a major overhaul is warranted. Having received feedback throughout this period, the policy owners may need to develop supporting standards to better assist control implementers to comply with the policy. This can be less complex than a high-level policy approval process, involving some supporting documents, e.g., standards or guidelines.
- **Retire**—Policies can be durable, but ultimately they all have a shelf life. Knowing when and how to retire policies is the job of the risk leadership. Sometimes it is self-evident when a policy should be retired; perhaps the organization has evolved, or regulations have changed. A new policy may arrive, subsuming an older one. When organizing controls across the policy framework, relevant sets of controls relating to a key risk may sometimes be bunched together. Retiring a policy goes through an approval process, and there may be a sunset period (e.g., 12 or 24 months) of agreement not to implement new controls, and not to subject new projects to enforcement of a policy that is going to be retired.

Tip: The policy approval board has a mandate from executive leadership, and that is normally reflected in the enterprise risk management framework.

Define Responsible Parties for Policy Enforcement

Typically, a policy may need sign-off by senior-level decision makers, whereas a standard may require sign-off only from the relevant stakeholder, such as the CIO, human resources (HR) or the risk team. Guidelines (e.g., on database security or cloud security) tend to fall within the purview of technology teams, with review by the risk team.

The organization should evaluate the cloud policy (or other policies applicable to cloud) in the context of the overall ERMF and technology strategy. Does the policy support and enable the needs of the technology strategy, or does it inhibit it?

To drive accountability, each control within a policy may have a designated control owner who is responsible for making sure the control is implemented and that it operates as defined. The control owner is accountable for making sure the control is enforced. The control owner may not be the designer of the control itself (the function may be carried out by a domain-specific or expert group). It may be helpful to think in terms of the RACI model: A control owner may be informed about the health of the control—any defects or noncompliance.

The shared responsibility model applies to cloud computing. The target organization will not have built all the controls, but it should have evaluated whether a user-facing control provided by the CSP is sufficient to meet the organization's control objective (e.g., need to know).

Some controls may be technical and others may be procedural; there may be an in-house mechanism for approving access requests for a particular cloud service, for example.

Note: This section covers what is considered cloud-specific. There may be elements that are not cloud-specific but that are still the responsibility of the organization to evaluate.

During the policy development stage, large organizations with multiple business units may have competing and conflicting cloud objectives or usage models. The cloud service user needs to determine how the policy owner has reconciled potential differences among different cloud objectives to arrive at a group policy that enables the organization to operate within its risk appetite, while fully leveraging various cloud services in the way the different business units wish to use them.

In the organization decision-making process, if there is a group-level architecture board or committee, cloud policy decisions will be influenced by that group's decisions. The group will record meeting minutes reflecting decisions taken. Part of good policy design allows room for stakeholders to agree to disagree, but there can be trade-offs and acceptances. Major stakeholders may agree or disagree with a policy. If they disagree, there may be a risk that they will act in bad faith or choose not to follow it.

Human Resources

HR is a key stakeholder and often an early adopter of cloud services, particularly SaaS. The cloud impact on policy is especially relevant, due to the way HR policies are designed. The primary requirement of HR is that records are kept, accessible and keyword-searchable, in keeping with reporting, alerting, and enforcement of acceptable usage guidelines. HR departments may need this assurance if they launch an investigation into an employee, or undertake an assessment of policy requirements concerning employee compliance with acceptable usage policies.

HR will want to know if the organization is using cloud services for collaboration. Using cloud email may also be relevant because of a CSP data retention policy, because HR may have reason to review emails as part of an investigation.

Privacy

Data sovereignty and data access rights are critical in a cloud context. A fundamental issue with cloud computing is the location where a CSP stores data. At a physical level, cloud storage can involve replicating data across multiple data centers, potentially in multiple jurisdictions. Of particular interest regarding data in cloud computing environments is the increasing commercial use and value of PII (personally identifiable information), and the sharing of PII across legal jurisdictions. The growing complexity of both ICT (information and communications technology) systems and data access requirements can make it difficult for an organization to ensure privacy and to achieve compliance with the various applicable laws. There is no single data access cloud control, since cloud services vary widely in their design. An organization may have policy requirements for third-party services to support role-based access control (RBAC). RBAC may be built into a cloud storage service, but it may require custom coding or integration with an existing data access solution.

A public cloud service provider is a PII processor when it processes PII for cloud customers. From a cloud architectural perspective, consideration often should be given to the locality of the organization's cloud data relative to the processor. Where is it going to end up in order to be processed? This creates legal and privacy challenges, which might include whether US and EU agreements on privacy are sufficient for the organization's needs. Large organizations will have privacy schedules that have specific privacy requirements that may be linked to country- or sector-specific legislation.

Procurement

From a procurement perspective, the organization should establish good cloud cost management practices. In part, good cloud cost management requires visibility into usage levels and procurement methods (e.g., employee use of company credit cards to purchase cloud services). It should centralize procurement around cloud billing, assessment of costs, and how well they match usage levels. Because cloud computing uses a consumption-based payment model, the organization should have billing alarms or other clear signals of overspending. A rising bill could indicate that an account has been misused or compromised.

The organization procurement policy should define mechanisms to identify misuse, waste or account takeover. It should apply best practices to billing and alert setup, setting soft or hard limits on spending proportional to what a given business function does. Is there evidence of good governance in place? Is the organization paying attention to billing trends?

Different environments will handle billing differently. There may be a playpen or sandbox for testing new cloud services, trialing ideas, and allowing teams to test the effects of different security settings. Signs to watch for are whether workloads are tagged as critical or can be removed with no impact to the organization. Are there barriers to prevent accidental or intentional overspending on virtual machines? The absence of those barriers are an unmitigated risk—it is the equivalent of giving an employee an unlimited company credit card.

Additional Cloud Policy Enforcement Guidance

Tip: It cannot be stated often enough: Just like the cloud journey itself, the organization cloud policy will evolve and be refined over time as the implementation matures.

Procedures and Technology

Procedures established under the cloud policy cover aspects such as user management. The target organization should have documented not just onboarding of cloud accounts, but the whole life cycle: how the accounts, master accounts and subaccounts are managed and monitored.

A technical control is one part of a broader set of mitigations, and the organization should consider not only how the control mitigates the risk profile, but also the life cycle of that control. Once the organization determines that it needs a control, it must define what it looks like, put it through a governance process, and then decide where to deploy it. If it is determined that the health of the control is good, the control should be subjected to conformance testing to validate that it is operating as it should.

(See sections 2.7, 2.8 and 2.9 for more details on controls and control objectives.)

Define Policy Violation Ramifications and Response Guidelines

The organization should have a means to identify, characterize and respond to policy violations. It should define clear and measurable targets and objectives, and use very clear language about the conditions for policy violations. In addition, policies and their objectives should be clearly communicated across the organization. In simple terms, users should know what they are permitted to do with the cloud service, and if there is doubt, how they can consult the policy and determine which procedure to follow.

Good cyberrisk management in all cases requires reducing the blast radius of a policy violation. Cloud services help in this regard, because the controls that a CSP gives to a cloud user can often be configured to limit damage from a policy violation. This is the secure by default approach.

The cloud customer, auditor or assessor should become familiar with the limits or step-up procedures that require the CSP to take action in order for a significant cloud event to take place, e.g., a deletion of a cloud account. Master accounts and subscription owner accounts are special and may require out-of-band interaction with the provider to delete those accounts or destroy fundamental data.

There are different ways to manage ramifications, from reducing the materiality of a policy violation to executing follow-up HR processes. Note that these actions are not likely to be cloud-specific. A good example is data loss leakage: It is necessary to assess the organization data leakage program to identify whether its controls also apply to cloud environments.

Response Guidelines

Many target organization response guidelines are not cloud-specific. There should be prudent risk management, with evidence of scenario planning for violations involving cloud services. In larger or more formal environments, that evidence would be in the form of meeting minutes detailing scenarios that were constructed and analyzed.

In smaller organizations, the guideline may be a one-page document outlining the scenarios it considered in terms of cloud risk events.

The response guidelines may depend on who is violating the policy and the access the violating party has—whether it is a full-time internal employee, a contractor or consultant.

Response levels vary. At the manager level, there is visibility into the cloud actions a particular user has taken, and the manager receives alerts for violations. This activity flows into standard company response processes. A violation might require escalation to HR for misdemeanors or compliance issues. Actions by users with greater responsibility, such as cloud administrators, represent a bigger material risk, so the response to any violations should be part of the scenario planning.

1.4.13 Cloud Governance Tools: Cloud Contracts

Cloud Service Agreement and Related Documents

The legal terms of cloud service agreements are frequently set forth in several documents, such as agreements, policies, rules, etc. For example, the following are commonly found in a CSP portfolio of legal terms:

- Master subscription agreement or master services agreement
- Data processing addendum (e.g., to accommodate the unique requirements of a specific law, such as GDPR)
- Terms and conditions
- Service level agreement (SLA)
- Acceptable use policy
- License agreement
- End user license agreement
- Services agreement (e.g., support services)
- Professional services agreement
- Website terms of service

There might be additional terms specific to a particular product or offering. Thus, the number of documents that may apply to a particular CSP service, their titles and content vary from one CSP to another. It is important to ensure that the CSP and the customer agree on which documents (and which version of those documents) constitute the entire agreement for the specific service being purchased.

Why So Many Components?

Although each potential customer may prefer having a single all-encompassing contract to govern the services it receives from a CSP, most CSPs use dozens of different forms of contracts to accommodate the variety of their offerings and the different choices each customer may want to incorporate.

In addition, the singling out of a particular concept in its own document may aid communications. For example, it may be easier to review an eight-page SLA to find the terms that pertain to the speed of communication, rather than to sift through a voluminous 88-page all-encompassing cloud service agreement that may not have the proper section header, forcing the reader to conduct a cumbersome word search to find the relevant section that addresses the topic.

One Contract Comprising Several Documents

The fact that there are several components does not mean there are separate, independent contracts. Typically, each document contains a provision that indicates it is part of a larger document.

For example, look for sentences such as this: “This XXX Agreement is made part of, and included in the YYY Agreement”; or “This XXX Agreement is made of the following Terms and Conditions, and also includes any document that refers to this XXX Agreement.”

These provisions are important because they are likely to affect the termination of a contract. If all components are deemed integrated into a single contract, then termination or breach of one of the contractual components may be deemed termination of the entire contract.

Conversely, other forms of contract may contain a clause with opposite intent, such as “Termination of this XXX agreement will not be deemed termination of the YYY agreement.”

It is important that all parties to a contract for cloud services understand the interactions between the various components of the legal terms, and the consequences if any provision is violated.

A contract formed of numerous parts may have some drawbacks. It is not uncommon to find provisions that conflict with each other. For example, one of the documents might provide for automatic termination of the contract upon expiration of the term, while another component may provide for automatic extension of the term from year to year, except upon notice by one of the parties X days before the end of the then-current term. To avoid these—often inadvertent—conflicts, carefully written contracts include a specific clause that defines an order of precedence for the components of the contract.

Contracts of Adhesion and Executed Contracts

The law distinguishes contracts that are executed by both parties (manually or with a digital signature) from those that are not, which are deemed contracts of adhesion.

In the cloud ecosystem, most contracts are contracts of adhesion—that is, the customer never signs a document, manually or digitally, but, instead, expresses consent by using the service. For example, the first page of a proposed contract may contain a conspicuously posted provision stating, “You automatically consent to these terms by using or logging into the service to which they pertain.”

The validity of contracts of adhesion is occasionally questioned, and litigation may ensue. This may happen, for example, if terms were not conspicuously posted or were egregious, or if they were changed without proper notification to the customer.

In rare cases, entities that have significant leverage or bargaining power are able to press a CSP to negotiate the terms of the offering, and in such a case, the entire agreement—or parts of it—may be signed manually or digitally by both parties.

Contract Terms May Vary Over Time

One of the most important consequences of the use of contracts of adhesion in lieu of executed contracts is that the terms of the former may easily vary over time. The provisions of a contract of adhesion take effect simply through a customer’s use of the service after having been notified of the modifications.

Changes in contract terms may have significant consequences for a customer. It is very important for a cloud customer to stay informed of the current terms of the contract and to monitor changes to the terms. Some CSPs may send a notice of change, and some may also provide access to the proposed new terms and indicate the date at which the change becomes effective.

Examples

The content of a service agreement or customer agreement may vary from one CSP to another, and it is impossible to provide a checklist or analysis of the different clauses, as they may be located in different components of a CSP contract.

A customer agreement, for example, may require customers to observe and comply with terms of service (ToS), a separate document that may define user rights and responsibilities, data stewardship, privacy policy, opt-outs and arbitration. It may also call out any terms that are specific to each service offered.

The ToS may be supplemented by an acceptable use policy (AUP), a narrowly scoped document that restricts the ways in which the provided cloud services may be used. The document may also outline the actions a CSP may take if the customer fails to comply with the acceptable use policy. For example, a ToS may state that the customer agrees not to attempt to bypass or manipulate security controls related to shared infrastructure, or the offending account will be terminated.

Service level agreements (SLAs) are generally used to specify the level of service that a provider promises to give to its customers and the standards with which it agrees to comply. They may set out the metrics for measuring the service, along with any remedies, penalties or service credits if agreed service levels are not met.

To illustrate how these agreements may interact, a customer agreement may state that the customer is responsible for properly configuring and using the services, or otherwise taking appropriate action to secure, protect and back up its accounts and content to provide appropriate security. The terms of service for a cloud service may require customers to include certain disclosures in the privacy notices displayed to end users of their products and services, and an acceptable use policy may prohibit the customer from offering content on a CSP platform that is not in-line with the applicable privacy laws.

In general, these agreements are not negotiated with each customer. The CSP sets the terms and conditions and the customer either accepts them or finds another CSP. Occasionally, an entity may be able to negotiate carve-outs, damages, or changes in liabilities, if it has unique leverage.

Terms of Service (ToS) Specific to Cloud

Similar to other types of contracts, the details of the contract terms help shape the relationship, the nature of the services purchased, the respective responsibilities of the parties, the conditions for termination of the contract, the representations and warranties each makes, and the allocation of risk. In most cases, the customer is unable to obtain changes to contract terms. In some cases, the alternative can be to acquire services offered at a higher price—but it is highly unlikely that the terms will be negotiable.

It might be more productive for the entity to try to mitigate some of the potential risks by designing and developing systems, policies, procedures and other structures that help compensate for shortcomings, anticipate potential new risk, supplement limited functionality, and otherwise add to the generic structure offered by the cloud service, addressing the gaps between the needed services and the limitations of the CSP.

The person in charge of the cloud service evaluation may help identify the gaps and related need for corrective action, and work with the organization to determine what risk decision it made, how it was reached, and how it was communicated.

Terms Specific to the Cloud

At the most basic level, cloud contracts contain the same types of provisions as contracts for other services. However, some provisions may have more importance and more drastic consequences in the cloud environment than in other settings, and thus may deserve more scrutiny. Consider, for example, the following:

- **Availability of the services on a 24/7 schedule**—What does it mean for support downtime?
- **Location of the servers, or of the CSP personnel overseeing the servers**—What effect does it have on compliance with foreign laws, such as privacy laws or data localization laws?

- **Access to customer data**—To what extent should the CSP have access to customer data, and what should be the limits on the CSP use or re-use of the data?
- **Customer responsibilities**—What do these responsibilities mean in practice in a distributed environment? How can the customer implement them in its internal policies or procedures?
- **Use of third-party service providers**—Most likely the customer will have no ability to object to the CSP use of third parties or contract staff. However, some privacy laws prohibit the use of subcontractors or subservice providers without the customer's prior consent.
- **Disputes and overdue charges**—What does the dispute resolution clause mean in practice? What leverage is left to the customer when the CSP can easily unplug the service?
- **Termination for breach**—If the sole remedy in a contract in case of breach is termination, what are the practical consequences for the customer?
- **Price changes**—What track record does the CSP have with respect to price and rate changes? What is the effect on budgeting?
- **Transition and data portability**—Is there such a clause, does it provide the customer with sufficient time and assistance to move its data to another CSP?

Changes in Service, Features, Terms

Cloud services by nature are continually updated; consequently, the related contract terms are likely to change to adapt to the changes in service.

In general, cloud agreements require cloud customers to accept changes made by CSPs as they evolve their services. Those modifications can have drastic consequences for certain customers and must be monitored to ensure continuity of the customer's operations. For example, the retirement of a feature might break a key functionality on which the customer depends, or it might impact user-facing controls.

The assessor should determine whether the cloud customer has an effective process in place to track and assess the impact of CSP changes in a timely way. The CSP agreement should be checked to determine whether customers are obliged to upgrade to the latest version of the service or if they can continue using the older version. If the cloud agreement allows the CSP to sunset any service within a set time, the customer may need to factor this risk into its own change plans.

It may be useful for the auditor or assessor to discover what cloud services the entity may plan to consume in the future, to be able to give a forward-looking view. The cloud customer should ascertain that the people responsible for implementing new features are aware of what SLA applies to those elements, and know that they are responsible for interpreting the cloud service-specific SLA changes.

Termination Clauses

Termination language in CSP agreements is designed to provide a path for the CSP or its customers to end the agreement in a transparent and coordinated way. It should describe a termination process that either party can trigger.

There are different types of termination, and they may have different consequences. Following are the most frequent types of termination:

- **Termination for cause**—This termination is triggered by a bad act committed by one of the parties, e.g., the customer failing to meet its payment obligation, or the CSP repeatedly failing to meet its performance requirements. The termination is intended to punish the party at fault.
- **Termination for convenience**—This type of termination occurs at the initiative of one party and without the fault of the other party. For example, the customer may be involved in a sale or merger, and the entity resulting

from the sale or merger wishes to consolidate different cloud services into a single relationship with a single CSP.

- **Termination at the end of the term**—This termination occurs when the term of the contract expires, and at least one of the parties does not wish to renew the relationship.

A well drafted CSP agreement should carefully describe the events that may trigger termination and the activities that are triggered by the termination of the contract. The more details provided, the easier the process is likely to be.

One of the particularities of CSP agreements is that even though a contract is terminated according to its terms, several activities must occur to ensure that the customer may continue its operations—should it wish to do so—after the termination. In most cases, the cloud customer will want to transition the service to a different provider, and the CSP contract should clarify how the transition will occur.

The customer-friendly termination clause should define the steps that enable the customer to retrieve its assets stored in the cloud, e.g., its data and data about its own customers. Many CSP termination clauses also provide a period during which the CSP will retain the data on its platform for the customer to download it. After that time has elapsed, the customer is responsible for gaining the CSP's assurance that the data was indeed deprovisioned, and not simply made inaccessible to the customer.

Some CSP agreements may contain CSP-friendly clauses, e.g., granting the CSP certain rights, such as the right to delete or aggregate any cloud service-derived data that relates to customer use of cloud services. An entity planning to use a cloud service should evaluate the effect of any such provisions, and consider whether they are consistent with its business model and meet any relevant legal restrictions.

The level of data lock-in will determine the ease with which a CSP customer can retrieve its data. Interfaces to export data from CSPs vary significantly but in general are easiest from commodity IaaS and PaaS providers and hardest from SaaS providers.

An additional consideration with leaving a SaaS provider is how to identify and obtain "tenant metadata," such as user account structure and roles, permissions and sharing settings, or indexes and dashboard designs. A SaaS tenant data-export facility may not include much in the way of metadata, which can lead to slow and costly migrations for large organization.

Even if a cloud customer may have evaluated the termination provisions during the sourcing process, keep in mind that most cloud services agreements evolve over time, and the termination rules may change during the life of the agreement.

The cloud customer should fully understand the termination process. Because the termination can be triggered by either side in most cases, the prudent cloud customer should devise reasonable plans to retrieve its data within the standard time frames stated in the agreement and keep those plans up to date.

Cloud SLAs

CSPs use SLAs to express expected results to their customers—usually related to performance attributes of a cloud service and the metrics used to assess them.

A cloud SLA defines the parameters for delivery of the cloud service, and what happens when something falls outside those parameters. The customer and CSP refer to it in the event of an outage or degradation of service. Cloud SLAs are important to several stakeholders within a single organization:

- For technology operations, the SLA defines what they can expect from the cloud provider service. The technology operations group will typically manage any deviation from that service.

- For the sourcing and finance functions, the SLA is linked directly to financial spend and sourcing due diligence.
- For the governance or control functions, the SLA is considered the reference for how events will be interpreted, particularly for assessors and auditors.

The key components of an SLA are outlined in **figure 1.18**.

Figure 1.18—SLA Key Components

	Component	Description
1.	Statement of intent or objectives	Outlines the reasons for and overarching purpose of the SLA
2.	List of services to be provided	Describes the duties of the service provider and customer, and any means for conflict resolution
3.	Availability and uptime	Time period and frequency for which CSP services must be available to the customer; uptime percentage usually measured and reported monthly
4.	Performance standards	Standards or benchmarks a CSP is prepared to commit to and deliver against; used to measure actual CSP service-level performance to ensure standards are met.
5.	Response time	Minimum and maximum amount of time that a CSP will take to respond to a request; may specify the form of approved requests (e.g., raising a ticket, rather than a free-form email to a help desk)
6.	Resolution time	Minimum and maximum amount of time that a CSP commits to resolve a task or issue

What has been done for performance in SLAs can also be done for security. The standardization community has been working to build SLAs for cloud computing through the development of ISO/IEC 19086.

ISO/IEC 19086 defines three important concepts:

- **Metric**—Standard of measurement that defines the conditions and the rules for performing the measurement and understanding the results of a measurement
- **Cloud service level objective (SLO)**—Commitment a cloud service provider (ISO/IEC 17788:2014, 3.2.15) makes for a specific, quantitative characteristic of a cloud service (ISO/IEC 17788:2014, 3.2.8), where the value follows the interval scale or ratio scale
- **Cloud service qualitative objective (SQO)**—Commitment a cloud service provider (ISO/IEC 17788:2014, 3.2.15) makes for a specific, qualitative characteristic of a cloud service (ISO/IEC 17788:2014, 3.2.8), where the value follows the nominal scale or ordinal scale

Build/Evaluate Cloud SLAs

In most cases, CSPs are reluctant to allow a third party to evaluate their SLAs and performance. There are practical issues (disruption to the CSP's business), financial issues (making staff available for such audits instead of attending to other tasks), and legal and liability issues (e.g., giving a third party access to some of the CSP's most prized assets). CSPs generally prefer to use external, independent auditors or attestations of security or compliance as a substitute for customer onsite audits.

Defining Metrics and KPIs to Monitor Provider Performance

Key performance indicators and metrics provide a window for assessing whether a CSP is delivering the services the customer expects to receive, at the required level and speed. They are important because the customer may be paying significant financial sums and trusting the CSP with potentially material amounts of sensitive business data.

Consistent failure by the provider to fulfill the metrics and KPIs is often an early indicator of a more significant issue between the customer and the provider. From the customer's perspective, defining metrics and KPIs forces internal decision makers to reach a consensus on what is important, or what constitutes a service operating within or outside—or materially outside—its risk appetite.

In theory, metrics are numbers related to service performance or customer requirements that are easily comparable. The source of most metrics for the CSP is the SLA, and the target organization must define its own metrics relevant to IT, risk and control owners. It then must map those metrics to the CSP activity.

Actions Required for Missed and Continual Missing of CSP Contract Terms

Some CSP agreements may include an escalating set of corrective actions and supporting operating procedures the provider will take. From the customer perspective, there ideally should be a document that describes the trigger points for those corrective actions and the CSP operating procedure.

The trigger for action may be in response to the CSP identifying that it has missed metrics, in terms of quality or speed. Ideally, the CSP contract should include a definition of what constitutes missed KPIs.

An action can simply be communication. SLAs typically define communication points linked to escalations. For example, if an event goes outside of tolerance and the parties have established a structure for periodic meetings, the incident will be discussed during the next account manager meeting call. If the event materially exceeds tolerance, it may be mapped to a priority one ticket within the IT governance operations, and so on, escalated through layers of management.

In some cases, the cloud customer may be able to define action at a certain point to involve redirecting traffic from the affected service to a standby service or an alternate CSP, or to communicate to customers that the service may be unavailable. If the cloud customer is a regulated entity, it must have a stated plan for how it would migrate from one provider to another.

Contract Terms and Allocation of Responsibility

Like other contracts for services, CSP agreements allocate responsibilities between the CSP and the customer. Contracts typically aim to establish a reasonable allocation of risks among the parties, taking into account market conditions, competitive offerings, pricing, demand, each party's needs, and the like. In the end, each party agrees to take certain actions and assume certain responsibilities. However, these responsibilities and risks concern *only* the relationship between the CSP and the cloud customer. Unless specifically and clearly stated in the contract terms, the CSP does not agree to be responsible to the cloud customer's own customers. The CSP is responsible only for providing the service that it committed to provide, and it is only responsible to the cloud customer.

Unless contract terms state otherwise, the cloud customer remains fully responsible to its own clients, customers, patients or end-users (collectively "end-user"), for any incidents that might have occurred while using the CSP services or that might have been caused by the CSP. While the cloud customer might be able to seek recourse against the CSP for damages it incurred while using the cloud services, most CSP agreements are not intended to establish a direct line of responsibility to the customer's own clients, customers, patients or end-users. There is no privity of contract between the CSP and the end users, and most cloud contracts do not make end-users third party beneficiaries of the contract.

Privity of Contract

CSPs and their customers establish the terms of their relationship through a contract. The contract is usually between two specified parties—e.g., the CSP and the customer—but, occasionally, it can include several parties, e.g., the CSP and all the subsidiaries or affiliates in a group of companies. The contract only binds the parties listed in the contract as having signed (manually or digitally) or having expressed their consent otherwise (by using the contract). Unless it is clearly specified in the terms of the contract, the concept cannot bind an entity that did not sign or consent to it. This concept is known as privity of contract.

This situation can occur in a cloud context: For example, assume there is a contract between a customer and CSP (e.g., for cloud services), and another contract between the CSP and a subprocessor (e.g., for hotline services). Assume the subprocessor hotline service is down periodically, which causes the customer to lose business. The customer cannot make a claim against the subprocessor for breach of contract because there is no privity of contract between the customer and the subprocessor.

The customer cannot seek compensation from the subprocessor under a breach of contract theory because there is no contract between the customer and the subprocessor. To obtain compensation from the subprocessor for poor performance, the customer needs to find a cause of action other than breach of contract. Alternatively, the customer should complain to, and seek recourse from, the CSP, which then needs to reach out to its subprocessor for compensation for the damages the CSP paid the customer, if any, as a remedy for the deficiencies in the subprocessor service.

Third-Party Beneficiary

Continuing with the example, it would have been possible for the customer to act directly against the subprocessor, if the contract between the customer and the CSP included a clause under which the customer was deemed a third-party beneficiary of certain provisions of the contract between the CSP and the subprocessor.

Occasionally, a contract may contain a clause that makes an entity that is not a party to a contract, a third-party beneficiary of certain provisions of the contract. This is the case, for example, in Clause 3 of the EU Standard Contractual Clauses Controller to Processor (SCC-CtoP).[27] Clause 3 provides that if a processor/data importer and a controller/data exporter have agreed to the terms of the SCC-CtoP, an individual whose personal data is being collected by the controller and processed by the processor is a third-party beneficiary of certain enumerate provisions of the SCC-CtoP contract between the controller and the processor, against the processor, or against any subprocessor that agreed to be bound by the terms of the SCC-CtoP contract.

Use of Third-Party Service Providers

In practice, cloud services are seldom provided by a single CSP. Even if a cloud customer engages a first-tier CSP, it should keep in mind that the CSP is likely to subcontract some services to third parties (subprocessors). CSP agreements usually contain a provision that gives the CSP the right to engage third parties to act as subprocessors.

Those provisions put the cloud customer on notice that third parties may participate in the provision of the cloud services. That does not mean, however, that the cloud customer is granted the ability to interact directly with those subprocessors, or to file breach-of-contract claims against subprocessors for defective service.

[27] European Commission, "Standard Contractual Clauses (SCC)," https://ec.europa.eu/info/law/law-topic/data-protection/international-dimension-data-protection/standard-contractual-clauses-scc_en

That is because there is no privilege of contract between the cloud customer and subprocessor; and there likely is no specific third-party beneficiary provision granting special rights to the cloud customer.

Provisions Required by Law

Continuing the example of a customer engaging a CSP, which then engages a subprocessor for matters that relate to the customer, it is clear that—especially in the cloud context, where CSP contract terms rarely can be negotiated—a customer seldom has the ability or the leverage to dictate to a CSP whether to engage a subprocessor, or the terms of the contract between CSP and subprocessor.

To cover that situation, and to protect individuals who would have no control over a subprocessor handling their data, several laws and regulations require that contracts involving the processing of personal data or personal information contain specified provisions to limit the right of the prime data processor to subcontract services to third parties. This is the case, for example, with Article 28 of the EU General Data Protection Regulation (GDPR), and provisions of the HIPAA regulations in the US (in the case of business associate contracts), which require that contracts with certain subcontractors contain specified clauses.

It is important to identify the particular kind of data that might be at stake, because national or sectoral law may impose contract provisions with specified requirements.

What/Why Internal Controls May Be Needed

Internal controls are needed to ensure that the organization correctly applies its security and privacy policies to its cloud activities. Employees could accidentally or maliciously onboard unapproved third-party services that might expose the organization to material risk, including inappropriate control of personal information (PI). This can lead to security incidents, such as network-layer attacks, and consequent reputational damage.

Numerous data leaks have been traced to cloud storage systems that held sensitive data in an easily readable format. Storage permissions are a CSP client-operated control under the shared responsibility model. The fault in these instances lay with the customer for poor implementation by failing to apply its own internal data protection policies correctly.

The policy control could be a technical control, e.g., an independent cloud security monitoring service that queries the cloud service to identify non-compliant integrations or extensions.

Some controls cannot be easily digitized and are implemented as offline paper controls that rely on employees' goodwill and adherence. Large organizations tend to engage an independent third-party cloud service for configuration, monitoring and assurance. It also may continually check for inappropriate integrations, unapproved storage or other control violations.

How Contracts Impact Action in the Event of Breach

This section explores two types of breaches: a data or security breach, and a breach of the service or contractual agreements between a CSP and a customer.

Data Breaches

For almost two decades, legislators throughout the world have made efforts to develop laws that address the consequences of data breaches—i.e., breaches of security that affect personal data. California, in 2002, became the first state to enact a law requiring private entities and government agencies to disclose to their customers or constituents the occurrence of a breach of security affecting certain categories of personal data. The concept has spread throughout the United States and to a large part of the world.

Most developed countries have incorporated data breach notification requirements in their national laws. For example, the EU General Data Protection Regulation (GDPR), which took effect in May 2018, requires data controllers and data processors to issue notifications when a data breach (as defined in the applicable law) occurs. Data breach laws generally require that the entity responsible for custody of the data notify government agencies and/or affected individuals when certain specified categories of personal data are stolen, accessed or modified without proper permission.

The incident that triggers the disclosure obligation is often identified as a "data breach". These laws usually require entities that process data on behalf of others to notify the custodian of the data about a security incident that might be deemed a breach. The custodian can then mitigate the consequences of the breach and complete the required notifications to the government agencies and/or the affected individuals.

Data breach laws contain most or all of the following provisions:

- Definition of what constitutes a data breach
- Data controller obligation to notify the relevant government authority, depending on the nature of the breach
- Data controller obligation to notify the affected individuals, depending on the nature of the breach
- Data processor obligation to notify the relevant data controller of the occurrence of a data breach
- Time frames for sending breach notices
- Information to be provided to the persons or entities to whom the breach notices are to be sent

While each law has its specific provisions, and each country has its way of reacting to a breach (such as litigation, class action lawsuits, or investigation by the country regulatory agencies or data supervisory authority), the common element is that data breaches have become a source of significant financial, legal, commercial and reputational risk to many organizations. Provisions addressing activities in case of a data breach, the interaction of the CSP and customer, and the related compensation or damages deserve scrutiny.

The customer should review the contract for specific terms concerning breach notification and for commitments the CSP formally makes, if any, for handling the process. If an organization needs to meet certain regulatory requirements, it should make sure the agreement covers the CSP ability to help deliver on those requirements.

The extent to which the organization is locked in to one CSP—e.g., if the provider holds all of its data—has a material impact on the actions open to it in the event of a breach. If the cloud service affects the customer organization supply chain, the customer must evaluate what services it provides to its own customers, and the services a CSP provides to the customer. If the terms do not align, then the organization might be in breach of contract in the event of breach of security at the CSP.

A distinction needs to be made between the customer's controls leading to the breach and the CSP controls leading to it. For instance, did the customer configure the user-centric controls in a certain way that led to the breach? Or was the breach due to a CSP defect causing a control failure that led to a breach? There may be a circumstance in which the CSP suffers a breach and the customer is left waiting to find out if its data has been impacted. In the contract, a CSP may want to exercise caution regarding breach notification commitments it makes.

Note, as well, that there may be a difference between who caused the breach and who is responsible for the breach notifications. Most data breach notice laws make the data controller responsible for the notifications to those individuals with whom it has contact, whereas an entity acting in a data processor (or service providers) capacity, is responsible for notifying its customer (the data controller, in most cases), but has no obligation vis-a-vis the data subjects (with whom it has no direct contact). Thus, it is important that the contract for services properly identify the respective obligations of the parties in a manner that takes into account their legal obligations under applicable laws.

Tip: It may be explicitly written that there may be circumstances in which a CSP cannot notify customers of a breach because the breach is being investigated by law enforcement, and the CSP has been requested not to make public the occurrence of the breach.

It is good practice for organizations to perform breach scenario testing prior to signing contracts with CSPs to explore and assess how it might react to a breach of contract by the CSP. In addition to planning prior to an incident, this can include scenario role-playing, potential actions during a suspected breach, and post-breach activities.

Under the shared responsibility model, cloud clients give up elements of control and have less visibility into communication flows and security events than with an on-premises model. The shared responsibility model requires more cooperation regarding actions related to a data breach. It may be useful to check whether, and how, a CSP, such as a SaaS provider, has helped with breaches in the past.

When an organization cloud-enables its incident response, it must accept limitations on visibility and control, which makes it more dependent on the terms of the contract with the CSP. The client's expectations should align with the service agreement it has signed with the cloud provider.

The use of cloud services can enable organizations to take an image of a specific virtual instance quickly and analyze it rapidly. Some CSPs enable customers to sniff networks, which can be useful when analyzing an incident scenario.

Numerous companies have incident response plans (IRPs) in place to guide their personnel on how to react to a security incident that might occur on the company premises or in its systems, or that might affect one of its service providers. To ensure that all relevant personnel are able to perform in emergency mode in accordance with the IRP, companies organize periodic tabletop exercises to train their personnel in how to identify a security incident and determine whether it may constitute a data breach under applicable laws. Some contracts for cloud services that involve large customers may include provisions requiring tabletop exercises to prepare for a potential security breach involving data stored in a third party's cloud.

Organizations should review their cyber insurance policies, in light of their agreements with CSPs to identify contractual remedies, conflicts or gaps in coverage or responses.

Breach of Contract

This section is concerned with allegations of breach of contract, where one party suspects or claims that the other party has breached their agreement, potentially or actually causing harm. While CSP-specific contracts vary, there are some common themes. Some are specific to the cloud, and others are common practices in contracts generally. With any contract, either party may breach its contractual obligations. Following are some examples of contract breaches.

These examples assume that the relevant CSP contract contains specific provisions that prohibit certain actions, and that a court evaluating the dispute would determine that one of the described activities breached a specific provision of the contract.

- **Client**
  - **Failure to follow instructions**—The customer may (after having requested permission from the CSP) engage in security testing of its workloads hosted at the CSP, but in the process harm shared components of CSP infrastructure. This situation might be a breach of contract if the contract includes provisions that require the customer to act with prudence when conducting its activities and ensure that its activities are not detrimental to other tenants of the CSP.
  - **Inappropriate disclosure of exclusive contractual terms**—A customer may have secured favorable terms with the CSP—e.g., it may have secured increased financial liability from the CSP. This situation might be

deemed a breach if the contract contains a provision prohibiting the disclosure of contract terms or specifying that contract terms are confidential.

  - **Bad tenant behavior**—Unsociable or extremely heavy use of a CSP management API can impact CSP operations. In extreme cases it can impact platform availability. Mature CSPs tend to have well-developed API throttling controls, but less mature CSPs may not. Extremely heavy use might be deemed a breach of contract if the contract contains provisions that would be interpreted as prohibiting certain uses of the platform. For example, a code of conduct, rules of the road or code of ethics document incorporated into the contract might prohibit repetitive activities that are likely to disturb the operation of the service.

- **CSP**
  - **Inappropriate disclosure of client information**—May be an issue if the contract contains a provision prohibiting the disclosure of confidential information and the definition of confidential information includes customer information or terms that can be interpreted as meaning or including customer information, and if other provisions of the contract do not exclude certain aspects of the client information, e.g., information that would have been disclosed to third parties without an obligation of confidentiality.
  - **Inappropriate access to, or use of customer's workloads or data by a CSP employee**—CSP help-desk staff might access, copy or reuse customer data for personal purposes, for example, or to develop a competing business. This activity might violate a contract provision prohibiting the service provider and its subprocessors from using customer data for purposes other than providing the service to the customer.
  - **Inappropriate access to, or use of the customer's confidential proprietary assets**—The CSP may access the customer's methods and/or source code, business model or processes for reasons other than to perform the services defined by the contract (e.g., to increase market share). This activity might violate a confidentiality provision in the contract to the extent that the data accessed and reused are confidential information protected under that provision. Most contract confidentiality provisions identify the information, documents, customer lists and research-and-development plans that are to be protected, and for which any use or reuse for purposes other than the strict performance of the contract is prohibited.

1.4.14 Cloud Governance Tools: Security Assessments

This section describes cloud security assessments. Although a cloud assessment or audit may be part of a broader organizational assessment, this section focuses purely on evaluating the cloud-related people, policies, technologies and techniques.

Scope, Timing and Nature of Cloud Assessments

There can be many triggers for an evaluation or an audit. One may be needed for contractual reasons, or because the auditee is a regulated business. Another trigger may be that a venture capital firm is interested in purchasing a business or startup that uses cloud technology and wants to check for any potential downside risk before investing. A CSP may have already published a security white paper that describes its data center and platform controls, and the security standards it has adopted. Although this approach is generally welcome, prospective customers may have internal policies that require suppliers to provide evidence of an independent certification of management controls against a recognized security standard. Beyond this, some customers will have higher assurance requirements that require CSPs to seek an independent attestation of their control design and operating effectiveness. Knowing your stakeholder (i.e., who has requested the audit) and knowing their requirements should guide the type of audit to be performed.

Scope

Scoping is critical to any evaluation, but it is even more important in the case of cloud services. When analyzing the supply chain in the cloud and establishing governance, the scoping exercise is key to identifying the different parties involved and defining what parts of a cloud service are or are not in scope.

When scoping any kind of cloud assessment, the assessor typically looks not at an audit or assessment of the CSP but at the organization (auditee) use of a cloud service. Therefore, the true scope of the evaluation will be the customer-facing controls of that cloud service. The assessor or auditor is allowed to look at the objects and services belonging to the cloud tenant or owner's account. The auditor may also look at some other objects and services but will not have permission to test them.

The auditor's role is not just to assess a particular service but to assess the way the organization designed and developed it and, more broadly, whether the design and implementation support and align with the organization's strategic objectives. Some considerations:

- How does the service interplay or interact with other services (whether internal or external/cloud services)?
- What is the inventory of data flows across the service supply chain, and what are the information security requirements and likely volumes of flows?
- How is user management handled, and what is its approval process?
- How are decisions about accounts reflected in an audit log?

The auditor will consider questions like the following:

- Were these controls and their designs checked by a third party?
- Is it possible to review the conformance tests to see whether the tool or service meets the owner's control objectives?

The auditor needs to consider the shared responsibility model when evaluating whether each control is present and effective. For example, to assess security awareness, an auditor may review policies and standards, looking for evidence of management requiring that all employees receive regular awareness training, testing (e.g., phishing simulations) and performance feedback. In the context of the cloud, the organization must also consider this control through the lens of the shared responsibility model:

- Does the organization (cloud customer) have an awareness training program in place?
- Has the organization (cloud customer) checked that the CSP has an awareness training program in place?

In more general terms, the auditor must evaluate the design and the operating effectiveness of each control (i.e., is the control working as designed?) in respect to the following:

- What is the organization doing for the portion of the control it is directly responsible for?
- What is the organization doing for the portion of the control that is under the CSP responsibility, but that the organization is still accountable for?

Two-dimensional thinking is a common theme that must be grasped and applied correctly to think like a cloud auditor.

Taking this theme deeper, following are some questions designed to help organize an auditee's thinking and shape the auditor's questions:

- Is it reasonable to expect a given control to be mirrored at both the CSP and the client? In the security awareness example, in which a CSP client is the auditee, both organizations would be expected to operate such a control independently of each other. The control objective and design may be very similar at a high level, but the implementation and operation of the control may differ. The CSP may be focused on reducing social engineering of its help-desk staff, for example, whereas the client (e.g., a bank) may be training its digital teams to recognize Internet-facilitated fraudsters. In the context of cybersecurity and the shared responsibility model, the client and CSP have similar control responsibilities, and both must be operated effectively to defend against social engineering attacks.

- Is it reasonable to expect a given control to be layered across the CSP and client? For example, the CSP provides its clients with logical access controls to manage access to cloud storage. This user-facing control is underpinned by both logical and physical access controls that the CSP operates—i.e., it implements them in the cloud platform. In this example, it is clear the overall responsibility for access control is divided between the CSP and its client. However, it also follows that even if the CSP platform is secure and operates within the CSP's stated control parameters, the CSP client can suffer a breach for failing to operate the user-facing controls properly.
- Is a given control operated by entities beyond the CSP and client? In complex multi-cloud/party cloud constellations, the auditor will need to consider how many hands are on the controls to identify the control operators involved. Taking this idea further, the auditor should then consider how many fingers of each hand are on the control to assess the materiality, or weighting, of each control operator, to help direct and estimate the level of effort and analysis the auditor should apply.

In summary, the role of the auditor is to take the following steps for each control:

- Identify the nature and placement of in-scope controls, along with control operators.
- Verify the auditee is operating the control within its organization (within the scope of the audit).
- Verify the auditee is asking the right questions of the CSP and other control operators.

Some cloud providers permit hands-on security testing of workloads and services operating on their platforms, but there are constraints and this is subject to permission. An IaaS provider may allow hands-on security testing of a virtual machine that is operated by the audit sponsor or account owner. However, the same testing on the cloud authentication service used by all cloud services would be out of scope.

Timing

Assessment execution options range from periodic checks to continuous assessment. The choice of option depends on the value of data and the risk associated with it. Traditional assessments tend to be annual and may be driven by contractual agreements. However, the cloud provides new ways of auditing, and the technology available to help with this process enables more frequent checking. This can lead to a set of hybrid assessments and audits (deeper but less frequent audits) or continuous assessments that are more frequent and automated, and that check for control and governance metrics.

The timing depends on the receiver of the findings. It may depend on the availability of key staff—or dependencies on others if the auditee is part of a complex supply chain—and on the availability of staff at subcontractors or at another provider. There may also be periods during the working day when assessment work is not possible—e.g., on a trading floor, an assessment activity cannot take place during regular business hours because of the potential operational risk.

Nature

Most cloud security assessments are logical in nature rather than physical—i.e., they involve checking software controls as opposed to physical security. Some cloud providers may not let the auditor of a customer visit their data centers, but at a minimum they should be able to give documentary evidence of the existing and emerging security standards they follow.

For the purposes of an audit, the auditor will need to treat some of the provider's security controls as a "black box," effectively, but will be able to query the user-facing controls that the customer can configure in order to implement its own security policies.

Some CSPs have developed and documented a shared responsibility model: The service provider is responsible for the security of its service, and the customer (auditee) is responsible for the security of the data. The role of the auditor is to assess the controls the cloud provider exposes to its clients and how the client uses them. It can be helpful to apply a RACI model to understand who in the auditee's organization is responsible for these controls, who is accountable, who contributed and who is informed. The auditor needs to establish the nature of the relationship between the cloud provider and its customer, which will determine exactly where responsibilities lie.

An audit can take two forms: a top-down risk based approach that starts with reviewing policies and standards to select key controls to audit; or a bottom-up process based approach to test the security of individual controls, e.g., a vulnerability scan that identifies gaps in security controls (see section 1.4.14). Regardless of the approach taken, risk will need to be evaluated and prioritized in a consistent manner.

Tip: From the practitioner's point of view, account design is more critical than ever when setting up a cloud environment. In a large organization, there is a master account and then suborganizations that create cloud user accounts, often given to divisional technology leaders. The account design model is driven by billing and security.

Types of Security Assessment

This section describes, at a high level, some common types of security assessments. It is up to the auditor to evaluate the capability and maturity of the auditee's security and to judge whether the controls are effective.

Audit Assessments

It is important to point out the distinctions between an audit and an assessment. The terms are sometimes used terms interchangeably, but there are crucial differences in the level of formality, assurance and admissibility as certain types of evidence. Audits can be internal or external but must involve a trained and certified auditor, whereas an assessment may be carried out by the organization or an auditor. An audit involves much more planning activity because the level of rigor is much higher than for an assessment. An assessment may employ some of the same principles, practices and processes as an audit—but it does not have to. Assessments can result in ambiguous findings; an audit should provide greater clarity. For example, an assessment might determine how many invoices an organization did not pay within 30 days; an audit would involve inspecting some invoices and checking specific payments.

External audit assessments can be carried out to check how a provider is upholding its end of the shared responsibility model. Some audit assessments are internal and conducted to identify cloud customer controls that are noncompliant with the organization policies or that fail to meet best practice standards.

Vulnerability Assessment

A vulnerability assessment is one part of the vulnerability management process. It is typically used in environments with considerable IT assets. A vulnerability assessment identifies and assesses technical vulnerabilities in an IT environment. It uses scalable commodity technology to carry out an automated series of technical tests quickly and cost-efficiently.

The result is a detailed technical report with a set of findings. An assessor should then review and evaluate the findings to eliminate false positives. This typically involves executing manual follow-up technical tests or fact-checking with the relevant IT administrators. A more challenging issue is how to handle potential false negatives—vulnerabilities that are present in the target but are not identified. Credible assessors will understand the appropriateness, capabilities and limits of their chosen assessment tools and techniques, and will consequently characterize any specific limitations, restrictions or caveats in their assessment report.

The next step is to assess, and potentially adjust, the default risk rating that was automatically assigned to the vulnerabilities identified by the assessment tools. This may be necessary to take into account the organization policies, risk rating definitions, asset value and ease of exploitation. In the context of IaaS, a traditional vulnerability assessment will scan customer-managed virtual machines to identify vulnerabilities in both operating system and application infrastructure contained within the customer's instances. Although this continues to be a best practice, it is insufficient because it does not include an assessment of the broader set of CSP platform user-facing controls. This leads us to a consideration of the emerging practice of cloud platform security assessments.

Note that cloud providers differ regarding notification requirements for vulnerability scanning, even if scanning customers' own servers on IaaS. Always consult the trust and security parts of their websites.

Cloud Platform Security Assessment

A cloud platform security assessment seeks to identify vulnerabilities or weaknesses in user-facing cloud controls. It is a review of the customer configuration of a cloud tenant. An example of an activity included in a cloud platform security assessment is reviewing whether the permissions on storage buckets are appropriate for the information they contain.

The scope of a cloud platform security assessment depends on the CSPs and cloud services that support and deliver the technology target of the audit. A complex or multi-CSP environment will require careful scoping to manage the size of the review. In general, the assessor should start with targeting a broad scope that encompasses not just the services used to deliver a particular business system, but all relevant and supporting cloud services. After sizing the scale of the resulting audit work, elements may need to be deprioritized or removed (with rationales clearly documented).

A simpler audit scope may be a lift-and-shift cloud hosted three-tier website. This may consist of one or more front-end web servers and application servers, and a backend database layer. An assessment of that environment would include both the contents of the virtual machines and the cloud management controls—such as cloud accounts, authentication settings, account privileges, access controls to the CSP management interfaces, virtual network configurations, instances of firewall settings, cloud storage security settings, etc.

More complex environments may include multiple cloud tenant subscriptions at one or more CSPs. For example, a software development company may use dedicated subscriptions per client, supplemented with multiple shared subscriptions for management accounts and shared cloud objects. The company may use a separate CSP for long-term data storage, so consider which elements to include. These scoping decisions are part of planning and should be made prior to commencing the review. The resulting scope should be challenged to ensure it makes logical sense and does not exclude key elements.

For SaaS environments, a cloud platform security assessor may be required to have business domain-specific knowledge (or be assigned a business partner for the review) to properly assess risk. It is one thing to identify technical security issues, such as single-factor authentication for highly privileged accounts, but another to identify that a particular set of roles assigned to a single account in a banking back end may lead to toxic risk combinations.

Penetration Testing

A penetration test, also known as a pen test, is typically a highly focused technical test of a technology, platform or organization. It is usually a manual or human-driven test that relies on the expertise, judgment and experience of the penetration testing team. Penetration tests may begin with a vulnerability scan, but unlike vulnerability scans, testers might try to exploit the vulnerability. Therefore, penetration testing should result in fewer false positives and better findings. One more difference between a vulnerability scan and pen testing is that scans are limited to the technology infrastructure and application, whereas pen testing can expand to include more techniques—such as human engineering, and manipulating the security processes in addition to the technology layer.

A penetration test generally results in fewer findings and fewer false positives than a vulnerability assessment. A subset of penetration testing is application security assessments, which attempt to identify business logic flaws and other technical weaknesses in custom business applications or platforms (as opposed to a technical device or off-the-shelf software). The test usually identifies the kind of vulnerability classes found in the OWASP Top 10.

Penetration tests vary in scope but typically are divided between application level or infrastructure level. In cloud implementations, penetration testing is primarily associated with assessing the application layer, which may itself be composed of multiple cloud services. In most cases, CSP customers are not permitted to test the underlying infrastructure. A growing number of security tools can automate some simpler aspects of penetration testing, but until there is a step change in tool intelligence, penetration testing remains an expert-driven manual activity.

CSPs with mature pen testing programs proactively test new platform features before they migrate to production. Some embed penetration testing techniques, tooling and personnel in their software development and build processes. Teams that test preproduction software early and often can discover and eliminate potential security flaws at a lower cost and with less operational risk than either a point-in-time penetration test or a vulnerability report from a security researcher or client. This shift-left approach being adopted as part of the DevOps approach reduces the cost to address vulnerabilities by finding them earlier in the development cycle. (See chapter 8 for more on this approach.) CSPs can deliver software features that their customers want faster and with measurable assurance. Consequently, point-in-time penetration tests tend to find fewer security vulnerabilities and become less of a security event. This arguably means they offer less value, and some seek to phase them out as a duplicate security cost. The counterargument is that a suitably scoped penetration test by a third-party expert provides independent evidence that software security controls are operating effectively in practice.

In light of these sweeping changes in software development, an auditor needs to recognize security assessment controls distributed across teams and processes. Examples include upfront threat modeling of a new software feature; mandatory security testing tollgates (which in turn, may underpin higher-level governance tollgates); and penetration tests of software features or modules. These tend to be smaller in scope and more frequently executed than traditional penetration tests, only validating new or changed software features but more often. Other leading-edge controls include the introduction of malicious software unit tests, e.g., SQL injection in login parameters, often codeveloped with penetration testers; security scans of source code for common vulnerabilities; and security analysis of third-party software components and dependencies.

With the rise of software feature penetration tests, an auditor should examine the decision making that drives the selection, scope, depth and timing of features that are tested. This includes examining whether the triggers for a test are appropriate given the role of the feature. A trigger could be a point in time, a new version of a software feature that implements a key security control, a regulatory requirement or client-driven demand.

In summary, the style of penetration testing as a security assessment is morphing, and mature CSPs are leading the change. A cloud auditor must go beyond a checkbox exercise of looking for evidence of a traditional annual penetration test by a CSP. Instead, the auditor should seek to identify and assess control elements embedded in the CSP software development and build pipeline, and qualitatively assess the decision making and application of penetration testing in practice.

Bug Bounty Programs

Some organizations offer bug bounty programs in an effort to draw on a broad and diverse pool of external talent to test the security of their services. These programs may augment—or in some cases replace—existing software assurance/penetration testing teams. Successful bug bounty hunters receive rewards based on the severity and quality of their security bug reports. When done right, bug bounty programs can increase confidence in the target service and provide a good return on investment. On the other hand, a poorly executed bug bounty program may attract few true experts and result in a deluge of substandard bug reports that an internal security team has to work through.

Red Teaming

Red teams challenge organization security by simulating an active and dynamic threat that seeks to achieve a well-defined and agreed-to threat objective. This objective may be composed of high-level goals to take control of or steal something of value to the organization. The existence of secondary goals underpinning the primary goals may allow the red team to demonstrate control defects even if primary objectives are not achieved. Red team engagements may be hands-on or tabletop exercises.

Unlike vulnerability testing and penetration tests, which are technology- or vulnerability-centric, red team assessments are threat-centric. Red teams execute adversary simulation exercises to assess an organization's cyber defense, detection and response capabilities. Exercise examples include breach of a credit card database, placing an unapproved trade that affects hedging, taking control of a utility provider, and stealing blueprints or other intellectual property. Full spectrum engagements may include social engineering of staff and bypassing of physical access controls.

Tabletop exercises are hands-off threat simulations led by an experienced facilitator or planner who identifies the relevant people (players) for each agreed-upon scenario simulation, which is subsequently played out around a table. The planner proceeds to inject hypothetical information to reflect unfolding events, e.g., members of the press just learned about the cloud breach and have started asking uncomfortable questions. The players brainstorm responses in real time and attempt to manage the response—possibly following existing response playbooks. The exercise continues with further injections and player reactions until the facilitator ends the role play and runs a feedback session with the group to stimulate learning and improvement. Follow-up tasks are assigned to implement learning lessons with the goal of becoming better prepared for the real event.

Cloud providers may place constraints on hands-on red team engagements, since cloud services by nature share infrastructure across multiple customers and an attack against one tenant could impact a shared component. In support of a primary goal, a secondary goal of a red team engagement may be to gain control of the cloud tenant master account. The red team may use technical or social means to steal credentials. In practice, the red team may only be permitted to log into the cloud interface with the master account, take a screenshot as proof, but go no further. These limitations should be set in advance during the red team scoping activity. CSP terms of service may disallow clients from executing red team testing, or make the tests nonviable through heavy constraints. Many large CSPs have their own in-house red teams that continuously test the effectiveness of their cloud platform and people. Other CSPs use third-party red teams. In both cases, CSP clients may request copies of test reports—but given the sensitivity surrounding such activities, they may have to settle for very high-level summary statements.

Purple team testing involves red and blue teams undertaking joint exercises to collaboratively test and improve key cybercontrols in a threat-informed way. Red team mimics adversary techniques, tactics and procedures (TTP), while blue team gathers valuable telemetry and observes the effectiveness of current or proposed defensive and detective controls/procedures. The primary benefit is that control defects—whether in protection, detection or response—can be quickly identified and remediated. A popular source of TTPs is the MITRE ATT&CK framework.[28] The MITRE ATT&CK Matrix for Enterprise applies to cloud-based TTPs.[29]

[28] The Mitre Corporation, ”Getting Started,” https://attack.mitre.org/resources/getting-started/

[29] The Mitre Corporation, “Cloud Matrix,” https://attack.mitre.org/matrices/enterprise/cloud/

Security Risk Assessment

Unlike other assessments, a security risk assessment is not characterized by a technical hands-on test of security controls. However, it can inform and be informed by such activities. As noted in section 1.4.10 on risk management, a risk assessment seeks to:

- **Identify the risk an organization may face through pursuing a particular course of action**—For example, adopting a new SaaS to support HR processes
- **Assess each identified risk to increase understanding of the risk**—For example, its source or driver; how the risk may manifest itself; the possible range, likelihood and frequency of realizing the risk impacts; its interplay with other known risk; availability and suitability of control options to diminish or mitigate the risk (e.g., the impact of a data breach or unavailability of the platform)
- **Respond to each assessed risk**—Considering context of the organization's risk appetite
- **Make recommendations to risk owners to accept a risk, develop compensating controls, transfer a risk or defer a risk**—For example, the SaaS platform may lack cloud data leakage controls, so the organization may decide to implement a cloud security proxy.

The output of a security risk assessment is a risk rating and one or more recommendations, which should be framed in terms of the organization policies, known control defects (e.g., in dependent controls), known plans and risk appetite. A recommendation can range from advice to proceed as planned, to proposals specifying one or more risk treatment actions to address identified risks. In some cases, the recommendation may be to defer or transfer risk (e.g., through cyberinsurance).

Ultimately, the risk assessment is intended to inform the organization risk owner's decision-making by providing a coherent view of the risks, along with options on how to respond to them. Risk assessments sometimes prompt a reevaluation of risk appetite by risk owners, or acceptance of risk. This can lead to a change of policy or creation of risk exceptions, which in turn may change the eventual risk rating and commentary. Management responses to risk assessments should be formally captured in the organization's risk management system. If it does not have a system, the assessor should provide an area for risk owners to respond directly on the risk assessment document itself. To ensure transparency and accountability, the assessor should recommend that the organization document any resulting policy changes, exceptions and control-related commitments.

Risk assessments vary in size and shape, according to the scope, materiality, complexity and ambiguity of the activity being assessed, the risk assessment methodology, the range of risks to consider, the number of risk scenarios to explore and controls to consider, and the number of stakeholders and experts to consult.

- **Scope**—An ambitious risk assessment may involve multiple organization technology platforms hosted at more than one CSP, crossing multiple business units in different time zones and crossing multiple physical borders. It is vital to recognize that the larger the assessment scope, the greater the risk to delivery.
- **Materiality**—A room booking system assessment likely has less materiality than that of an online banking platform. It follows that you should dedicate more time to assessments of greater materiality. The assessor needs to understand what is considered material to the client and its stakeholders. It is not unusual for a client to seek an assessor's opinion on where to focus, which ultimately is a judgment call. A helpful question to ask of an accountable risk owner is: What cloud risk keeps you awake at night? A helpful answer will refocus attention on what is ultimately considered material. The materiality can help drive the choice of assessment type—e.g., cloud platform security assessment (good coverage, average assurance) versus a penetration test (focused, expert-driven technology assessment) or red team engagement (adversary simulation for critical business functions).
- **Complexity**—Discussions in the planning process should focus on identifying and eliminating unnecessary scope complexity to balance the number of moving parts to be assessed with the goal and ultimate value of the exercise. Assessments may be broken into phases to make them more manageable, or segmented and assigned to

different assessment teams—e.g., a SaaS HR domain expert has a narrow role in a broader organization cloud review.

- **Ambiguity**—Elements of ambiguity or uncertainty, such as a pending choice of CSP, must be highlighted and resolved. In some cases, the simple answer is to defer the assessment until there is greater clarity. In other cases, the assessment is required to inform a decision that will resolve the ambiguity—e.g., comparing the implementation of two competing cloud designs, or assessing the effectiveness of different CSP DDoS controls. A clear scoping statement addresses both the source of ambiguity and how it will be handled.
- **Risk Assessment Methodology**—The choice of methodology and target level of adherence naturally drive the depth of activity and the choice of assessment tools. It is important not to confuse a risk assessment methodology with either a risk assessment standard or assessment tool-specific methodologies—i.e., a professional penetration tester usually follows some sort of methodology to execute the test, but this methodology is distinct from the overall risk assessment methodology of which the penetration test may be just one component.
- **Range of risk**—Large organizations with mature risk capabilities break out their risk management practices into principal or key risks, with separate teams per risk area. A cloud environment may be assessed for cyberrisk, information risk, technology risk, privacy risk, compliance risk, sourcing risk, etc. There may be organization-specific risk policies, measures, metrics and thresholds to review and use, and they may themselves need to be mapped to the risk assessment methodology. Some organization policies may mandate a risk assessment for large changes, which may require penetration tests for completion prior to going live.
- **Number of risk scenarios**—A risk assessment may seek to evaluate a range of cloud-related events that can lead to a business impact, e.g., a cloud storage service outage, a CSP data breach, or a CSP going out of business. Scenario-based assessments lean more toward red team engagements, e.g., a tabletop exercise with relevant personnel taking part in a simulated exercise to explore a given risk scenario.
- **Controls under consideration**—The number, location and complexity of controls often directly inform the choice of assessment tool. A SaaS provider seeking a technical risk assessment of its platform may opt for an application security assessment and infrastructure penetration test. One of the SaaS providers customers, however, would likely choose a cloud platform risk assessment to review its user-facing controls of the same SaaS. Taking this example further, the SaaS provider may host its platform on a third-party IaaS. A large prospective client may identify this subcontracting relationship during its own preboarding risk assessment, and as a condition of signing up to the service, may require the SaaS provider to undergo an independent cloud platform risk assessment of its use of the IaaS during the subsequent 12 months.
- **Number of stakeholders**—More stakeholders means more meetings. Effective stakeholder management requires a clear understanding of stakeholder needs (especially frequency and style of communications) and stakeholder expectations. It is easy for stakeholder engagement to pull time from assessment activity, which can reduce the effectiveness of the assessment. Cloud assessments in particular may involve multiple business partners, multiple CSPs and multiple internal managers.
- **Availability of experts**: The cloud technology landscape is vast, and new types of cloud services are announced frequently. Cloud security experts are in high demand, but they cannot possibly have deep expertise across all services. Assessors need to be transparent about skill shortages to avoid missing key issues and giving a false sense of security. Third-party assessors may augment internal assessment staff, particularly to assess niche technologies (e.g., cloud HSM for key storage) or complex multi-cloud architectures. In addition, some organizations may place restrictions on the individuals who can perform risk assessments—e.g., the ministry of a nation state may only permit country nationals. Also important is the availability of experts at the target organization who may be fielding requests from many customers. It is vital to identify which experts are needed, and to structure a risk assessment and secure the required resources. If relevant experts are not available, it may be necessary to weigh different potential outcomes or devise scenarios that assume certain conditions to be true.

There is a wide variety of security assessment tools available for different needs and use cases. A key element of risk assessment planning is understanding which tool is appropriate for a given context. Simple risk assessments may only require the use of one tool, whereas others may employ multiple assessment tools across multiple phases with each phase informed by the results of the prior one.

The "Cloud Octagon Model" is a resource that can help organizations with assessing SaaS providers.[30]

Limitations

Cloud services are built on the foundations of other services, so the shared responsibility model places limits on the kinds of assessments an organization can carry out. Most cloud providers do not permit penetration tests or vulnerability scans of shared infrastructure, but large CSPs publish assurance statements that describe their controls and how they are evaluated, and a security testing policy that often evolves over time (to allow more testing with less advance notification). The policies usually allow more testing as they mature.

The assessor will mostly rely on a cloud provider statements. If a provider offers a summary of the results of the assessments carried out by a third party, the assessor will want to know if the scope of those assessments is material and sufficient to the services that the auditee is consuming. The scope is critical, because it refers to attestations of security. Effectively, the auditor wants to establish the relevance, quality and depth of any third-party security assessments: Who performed the assessment (testers' skills, expertise and credentials)? How much time did they spend on the test? Were they operating under any test constraints?

The assessor may engage with the cloud provider security team, and the nature and speed of the team's response to security questions could form part of the auditor's qualitative judgment. The auditor may also establish whether the person responsible for security at the auditee simply accepted evidence of security testing from the cloud provider or asked follow-up questions to establish more clarity around security controls.

1.4.15 Cloud Governance Tools: Auditing Overview

Audit reports generally have a higher level of formality and credence compared to internally produced assessment reports. First, establish the criteria for the planned audit activity: What is defined as success? Success as defined through audit reports can vary greatly by organization and by audit type. One organization may rate audits simply pass or fail; another organization may have ratings per individual finding, observation or exception, plus an overall rating for the area under review. When a regulatory entity performs an audit, it may define success differently because the bar for a compliance audit may be higher than for an operational audit.

See chapter 5 for additional details on cloud auditing.

Audit Types

First-Party (Internal) Audit

A first-party audit is a self-directed audit in which the auditor and the auditee are part of the same entity. In some sectors, such as banking, internal audit may be a standalone function or department, separate from the risk function, with a separate reporting line from operational leadership. The head of audit may report directly to the CEO and, in many cases, the board of directors (or the audit committee of the board of directors). Internal audits are formal and in support of independence should follow the auditing principle of objectivity. Objectivity is facilitated for internal auditors in that they are prohibited from auditing an area where they previously had managerial responsibility for a designated time (often referred to a cooling off period). The intent is for auditors to avoid being in a position to audit their own work. In practical terms, internal audits can carry significant weight, as reports are generally shared with

[30] Cloud Security Alliance, "Cloud Octagon Model," 24 June 2019, https://cloudsecurityalliance.org/artifacts/cloud-octagon-model/

the organization's regulators and external audit firms. Sharing of reports is part of the effort of internal and external auditors to collaborate and ultimately avoid duplication. External audit firms also share work with internal auditors. In such cases, the internal auditors will then work with the organization to monitor the external auditors' findings to resolution.

Tip: Due to the nature of cloud services, the organization loses a certain amount of control compared to an IT service delivered in-house. For this reason, it is of paramount importance to reinforce the role of the internal audit function, and to establish an auditor mindset and an accountability culture within the organization.

See chapter 5 for additional details on the role of internal audit.

Second-Party (External) Audit

In a second-party audit, the auditor and the auditee are not part of the same entity. An external audit can be gauged by the level of independence and audit expertise that otherwise would not be available to the organization. For external audit firms that are subject to the Public Company Accounting Oversight Board (PCAOB), there is guidance on audit partner rotation to facilitate independence of the firm relative to the organization. There are different requirements, or triggers, for an external audit, and the scope may vary widely. The organization may want to use the audit result for marketing purposes, to satisfy a contractual obligation for a customer, or to fulfill a regulatory requirement.

Third-Party (Certification/Attestation) Audit

A third-party audit carries the most weight with external stakeholders and tends to have the highest level of acceptance given the level of formality, the robustness of the auditing methodology, and the strict competency requirements of the auditors. It is typically performed by accounting firms (e.g., SOC2, CSA STAR Attestation, BSI C5), regulators, or specialist audit and security companies that are formally accredited by national accreditation bodies, or industry bodies that set and enforce standards (e.g., ISO 27001, CSA STAR Certification, PCI DSS).

An attestation is a document prepared by a skilled independent practitioner, stating that the controls an organization has in place are effectively designed and implemented. On the other hand, a certification, which is issued by a certification body or industry body, states that technical and organizational control objectives or technical measures are effectively within the scope of the certification.

How to Audit the Cloud

This section introduces and provides an overview of how to audit the cloud.

Devise an Audit Plan

An audit plan is the set of activities that comprises the audit. The choice of how to audit an auditee plays a major role in determining how much time will be required and identifying dependencies. This is best decided up front, prior to agreeing on the audit timeline. The auditor will need to perform exploratory activities to identify which artifacts are available for review and what level of access to the target environment is practical (both physical/offline and logical/online), and gather availability and contact details of relevant subject matter experts. Sequencing of activities will depend primarily on completing dependency activities first, and by the timely availability of any target-provided artifacts and their subject matter experts (SMEs).

At a high level, the auditor is seeking to ascertain that the auditee is doing what they say they are doing—i.e., to ensure in-scope controls are in place and functioning as designed; not doing something that would materially impact what they are doing; and identifying material errors, conflicts of interest or omissions as they relate to the audit scope (and not beyond). For example, a client that claims to securely handle personally identifiable information (PII) may use a SaaS provider that employs encryption at rest for attachments uploaded to its service. In response to user complaints about limited upload quotas, the client's SaaS tenant administrator has installed a marketplace plugin that enables users to seamlessly store larger attachments—but at a third-party provider that does not encrypt them. The auditor needs to recognize this situation and identify if the client is uploading large attachments containing sensitive data, because that likely would breach data protection regulations.

The auditor should be able to answer the question, What am I auditing against?—i.e., the standard or documented set of controls used as a basis for determining the truth of an auditee's compliance.

Auditors should apply the relevancy test early in the audit process, as they develop an understanding of the audit environment, to better identify any controls that should not be audited—e.g., because they are irrelevant, or because the activity is in the planning phase and not yet formally defined and approved within the organization's control framework.

Note: With formal audits, the auditor is not playing the role of advisor, providing professional consulting, or teaching the auditee how to check the box or cheat the audit. In addition to carrying significant reputational risk for the auditor and its employer, such behavior is contrary to the standard auditor code of ethics. It can lead to disciplinary action, such as loss of employment, and/or criminal prosecution leading to fines and jail time. It is acceptable to share common approaches that other organizations may take to meet an audit requirement, but auditors should never identify the source or share anything proprietary or confidential from other organizations.

Auditing Security Requirements

The audit scope of security requirements goes beyond assessment of technical security controls. It encompasses examination of governance processes, risk management, cloud policies/standards and third-party assessments. The design of a cloud governance framework should consider the scope of cloud activity across an organization. From an audit perspective, an organization's cloud workloads can be classified as falling within its governance framework (assuming it has one) or outside. An example of the latter could be an unapproved SaaS source code repository used by developers for hosting the company source code (just one example of shadow IT).

A workload may be deemed out of governance because it is not formally identified and tracked within the auditee's governance framework, or because the organization has assessed it as noncompliant. It is important to make these distinctions when scoping an audit and expressing audit findings.

Governance processes flow from and reflect the auditee's implementation of its chosen governance framework. The auditor first needs to evaluate whether the reasons the auditee adopted its chosen governance framework continue to exist, and check that its interpretation of the framework, in the context of the scope of a cloud audit, is rational and reasonable.

The amount of audit effort involved in governance processes will depend on both the purpose and scope of the audit and the auditee's maturity. If the organization has only recently adopted the framework, it is unlikely there will be much to examine. The focus should then shift to evidence of proper planning, decision making and initial implementation.

Since governance frameworks can be large, the auditee may have a governance implementation plan that reflects the steps it is taking to adopt the framework and achieve compliance in different parts of its business. Such plans seek to establish or improve the maturity of governance processes that ultimately bring the governance framework to life.

The auditor should review plans that pertain to governance of the organization's cloud activity to calibrate expectations, assess maturity and refine the audit plan.

Credible plans leave a trail of evidence that the auditor can review. Senior leadership will approve and track the plan, and there will be line items in official budgets and headcount plans to deliver the required support. In some cases, plans may be set as performance goals for responsible managers.

Regardless of the scope of the audit, the auditor should become familiar with the auditee's governance processes, which provide the foundation for the organization's risk management practices, influence policies and standards, and set expectations for third-party assessments. This will not only help the auditor speak the language of the organization, but also navigate the lifecycle of governed cloud workloads, and therefore is vital to the success of the audit.

The cloud auditor will want to examine documentation supporting relevant governance processes for their underlying purposes, goals, inputs, execution steps, outputs and measures. Once the governance processes related to cloud operations are understood, the auditor can assess the extent to which they: are fit for purpose; reflect and align with the governance framework; are applied appropriately; improve over time in light of experience; and ultimately deliver security outcomes aligned with the organization risk appetite.

Tip: The presence of good governance does not guarantee good security practices. However, it is rare to observe the latter without evidence of the former.

With the governance practices understood, the auditor can turn to understanding and assessing the auditee's cloud risk management activities. The enterprise risk management framework (EMRF) provides a blueprint for how an organization identifies, measures and manages risk. It will typically outline the principal risk the organization faces, and describe the mechanisms in place to measure and manage those risks (including exceptions).

The detail and number of supporting documents will depend on the level of regulation in the auditee's sector and its overall maturity in dealing with risk. To keep things focused, risks associated with the cloud will typically fall under cyberrisk or technology risk, which themselves fall under operational risk. This is not a hard and fast rule, and the auditor should consult the auditee's ERMF as the definitive road map.

The auditor should ask the auditee to highlight the relevant parts of the framework, bearing in mind the goals and scope of the audit, and walk the auditor through the core risk management process as it relates to cloud workloads. It is not uncommon for auditors to risk-assess for the audit themselves. They may rely on the ERMF as a background or foundation but still tailor the risk assessment to fit the audit's objectives. The auditee should be able to substantiate its risk management practice with evidence the auditor should assess across each core risk management activity:

- **Risk identification**—Core practices include concern reporting, which encourages and enables staff to report potential risks in a consistent and timely manner. Reported concerns are then evaluated by suitably competent persons for relevance, materiality, likelihood and toxicity. Confirmed risks may be logged in a risk database and reported upstream to risk owners and risk committees (depending on materiality). Major risks may be recorded on an organization- or function-wide risk register. For example, a help desk operator may report that customers are calling in to report that when they log into the company website they are seeing other customers' details.
- **Risk measurement**—The auditee may have defined qualitative and quantitative risk metrics and measures, which may be compared against management agreed-upon values that reflect risk appetite and drive an overall risk status—e.g., operating within risk appetite, approaching risk appetite, exceeding risk appetite. Cloud-related risks may not be measured directly, but may fall within a broader category. In a conglomerate, risks will be aggregated across organizational layers, and a top-level risk committee may meet to determine whether any response is required. Minutes of these meetings will be recorded and will include packs (or decks) provided by relevant risk and control experts, which can provide valuable insight to the auditor.

- **Risk management**—Key tools for managing known cloud risks are policies and standards that describe what controls should be in place and, in some cases, how they should be implemented. Policy language varies from "principles-based" to "prescriptive." The auditor should read and assess these policies and standards carefully, and seek evidence of how the auditee achieves compliance. Exceptions and noncompliances may signal how the risk processes work in real life. Discrepancies should be investigated carefully and followed up diligently as they sometimes can reveal underlying behaviors or failures that are accidents waiting to happen.

Supply chain or third-party risk is an example of an adjacent risk that is often relevant to cloud risk and will typically be in-scope of a cloud audit. Key artifacts to examine are third-party risk assessments of CSPs. Some organizations will have policies that require on-site audits, which can present challenges when dealing with large CSPs. However, for smaller CSPs this requirement is more reasonable and the auditor should look for evidence that a third-party audit was performed.

Third-party audits are becoming increasingly standardized and follow a common pattern. In some cases, they are outsourced to generalist organizations that may lack cloud expertise. The auditor should assess the method, coverage and completeness of prior assessments and evaluate the depth and breadth of coverage. An emerging best practice is the move toward a continuous auditing approach that reduces the dwell time between point-in-time checks and increases their frequency to provide a higher degree of ongoing assurance.

Auditing Legal and Regulatory Requirements

Organizations are subject to legal and regulatory requirements, and an auditor must have a clear understanding of the jurisdictions in which the organization operates. CSPs typically do not make legally binding claims that their service is fit for use in a particular jurisdiction, even if they already operate in the same jurisdiction. Rather, a CSP confirms that it complies with local laws and any regulatory requirements in the jurisdictions in which it operates.

The cloud customer is ultimately responsible for establishing whether its own activities, including its use of a particular CSP, comply with local laws and regulatory requirements in the jurisdictions and sectors where it operates. Given the potential fines and sanctions for failure to comply, it is vital that the auditor correctly identify the set of applicable laws, requirements and compliance requirements to audit an organization against.

An auditor who has limited experience auditing legal and regulatory requirements of the type and nature of the audit target should make it known during the audit scoping phase and arrange suitable expert support. This arrangement could be co-sourcing or outsourcing. Co-sourcing has the advantage of the auditor being able to develop professional expertise in auditing legal and regulatory requirements. At a minimum, the auditor should review draft findings with a relevant expert. Failure to do this can lead to professional liability, which may not be covered or may have limited coverage by any professional indemnity provider.

Auditing SLAs/Performance

With the cloud, organizations do not have direct access to the CSP physical infrastructure and cannot test it. However, they can use data provided through SLAs, third-party audits and compliance reports based on laws and regulations.

When examining SLAs, an auditor ascertains whether the CSP SLA enables the control owners to operate a service at the stated level. Often, the auditor looks for what is missing in the SLA, such as key operational matters that are material to the operation of the service.

Challenges of First-party Audits in the Cloud and Suitable Alternatives

A customer that uses a public cloud no longer has direct control over several IT security and privacy-related features. CSPs typically prohibit security scanning, penetration testing and first-party audit by contract. If the CSP has audited itself but there is no external attestation, there is a high counterparty risk, which most organizations want to reduce.

However, with experience and growth, CSPs generally start providing higher assurance statements, either in the form of an internal audit or one conducted by an external security company. If cloud customers need more assurance—such as in regulated industries, or in the case of an organization requiring accreditation for contractual reasons—the CSP may engage a certified practitioner or firm to provide an attestation of compliance.

Another suitable alternative is a community cloud, which is a shared infrastructure for organizations, e.g., in the same sector or with similar needs. In a community cloud, an attestation of security benefits all users.

See chapter 5 for additional details about cloud auditing.

1.5 Chapter 1 Knowledge Check

REVIEW QUESTIONS

1. Which of the following elements are part of cloud policy? (Select all that apply.)

   A. Types of data that are allowed to migrate to cloud
   B. Relevant stakeholders
   C. Ability to perform service catalog
   D. Mandatory controls

2. When evaluating a cloud provider's maturity and ability to execute, the customer should consider (select all that apply):

   A. Fiscal performance
   B. Board involvement
   C. Transparency
   D. Number of mergers and acquisitions (M&A)

3. For cloud trust and transparency, what are some key considerations that the cloud customer needs to be aware of? (Choose two.)

   A. Inadequate transparency and auditability by CSPs
   B. Lack of adequate cloud services and resource testing
   C. Loss of governance based on distributed services
   D. CSP service contracts address all client security risk in an IaaS deployment model

4. Considering that cloud security is based on the shared responsibility model, select the controls that are usually the responsibility of the cloud consumer:

   A. Granting or removing application user access
   B. External cloud backups
   C. Security of the hosts and hypervisor layers
   D. Producing security scanning reports

5. In SaaS, which risk has a shared responsibility by both the cloud customer and cloud security provider?

   A. Infrastructure-related risk
   B. Application-related risk
   C. Data classification-related risk
   D. Identity and access-related risk

6. Which of the following parties should be included in a cloud service provider (CSP) compliance scope? (Select all that apply.)

   A. Subcontractors
   B. Subcloud service providers
   C. Employees
   D. Cloud customers

7. When reviewing proposed policies, what is the remit of a policy approval board? (Select all that apply.)

   A. Whether the implementation period is realistic
   B. Consider the potential cost implications of implementing the policy
   C. Determine the fit within the existing policy framework
   D. Determine cloud technology strategy

8. Which of the following are different types of hands-on security assessment?

   A. Vulnerability assessment
   B. Cloud platform security assessment
   C. Security risk assessment
   D. Tabletop exercise

9. What are standard audit types? (Select all that apply.)

   A. Internal
   B. Self-assessment
   C. External
   D. Third party

10. Match the requirement to the correct requirements scheme.

Requirement	Requirements Scheme
A. Define SLA requirements	1. Security requirements
B. Understand the applicable law	2. Operational requirements
C. Define controls framework	3. Legal and regulatory requirements

11. Match the terms and their definitions.

Term	Definition
1. Termination for cause	A. Used when circumstances have changed, and a party needs time to reevaluate its options.
2. Termination for convenience	B. Used when circumstances have changed, and a party is able to stop without having to pay damages or penalty.
3. Suspension	C. Used to identify a contract that continues for a longer duration than originally expected.
4. Cure period	D. Used to identify a potential consequence of bad behavior on the part of one party.
5. Extension	E. Used when a party who has (or is accused of having) violated the contract is given an opportunity to correct the bad act.

Answers on page 102

Chapter 1 ANSWER KEY

Correct answers appear in **bold** font.

1. **A. Part of cloud policy is defining controls which are mandatory.**

   B. Mapping relevant stakeholder is part of cloud policy.

   C. Cloud policy is about security and resilience requirements.

   D. Part of cloud policy is defining controls which are mandatory.

2. **A. Market share, revenues, budgets and other fiscal meters are all indicators for maturity and ability to execute (unit 1.3.1).**

   B. Senior management and board involvement in security strategy is in indication for how significance is security to the organization and an important sign for ability to execute (unit 1.3.1)

   C. Organizations who deploy mature information security policies tend to be more transparent since they are confident in their decisions and able to clearly defend and explain their strategy. This is indication for is maturity (unit 1.3.1).

   D. Mergers and acquisitions are not an important factor to consider when evaluating the provider maturity because while M&A could be important to acquire new capabilities or build new synergies, do not necessarily have a positive impact on the level of maturity of the service and of the CSP.

3. **A. As time has progressed CSP tools and capabilities facilitating the auditability of both CSP and cloud customer environments have become significantly more effective.**

   B. Similarly to question 1 CSP capabilities in this area have become more numerous and can more effectively benchmark and visualize the testing of cloud native tools, applications and workloads, this answer is incorrect.

   C. Due to the abstracted and decoupled nature of many of the CSP services and resources it is a key responsibility of the data owner to understand the jurisdictional implications of running services in different regions (General Data Protection Regulation or California Consumer Protection Act considerations would be an example of potentially 'imposing' governance outside of their respective jurisdictions). Organizational and local data protection regulations must always be at the forefront of any cloud strategy. This answer is correct.

   D. This is not possible, due to how risks evolve within a cyber-security context. In addition, within the Shared Responsibility Model, the provider is only responsible for the areas laid out in the service contract - furthermore providers will naturally look to reduce their risk profile as much as they can in any contractual situation. The customer must be aware of this, consequently this answer is incorrect.

4. **A. Paragraph 1.5.4. Examples of controls that are usually in the responsibility of the cloud consumer: access controls.**

   B. Paragraph 1.5.4. Examples of controls that are usually in the responsibility of the cloud consumer: External cloud backups.

   C. Paragraph 1.5.4 Examples of controls that are usually in the responsibility of the cloud **provider**: Security of the hosts and hypervisor layers.

   D. Paragraph 1.5.5. A number of security services and capabilities, like for instance, **security scanning**, penetration testing and first party audit are typically **prohibited** by contract.

5. A. Incorrect (cloud provider only)
   B. Incorrect (cloud provider only)
   C. Incorrect (cloud customer only)
   D. Correct (section 1.4.10). Based on Cloud Customer and Provider Shared Responsibility Risk Matrix.

6. **A. Should be part of compliance efforts**
   B. Are part of compliance efforts (e.g. IaaS cloud provider used by a SaaS provider)
   C. Are part of compliance efforts
   D. As per (Unit 8, 1.3) cloud customers are not part of compliance efforts.

7. **A. The board will balance the demands of other policies in implementation periods with major upcoming business initiatives.**
   B. The board will weigh up cost implications of a proposed policy and may demand/challenge cost estimates.
   C. The approval board will consider alignment with and duplicative/overlapping controls in existing/proposed policies, level of detail, tone, clarity, style and terminology.
   D. This is not within the remit of a Policy Approval Board.

8. **A. The assessor uses vulnerability scanning tools to probe systems and services to identify technical vulnerabilities.**
   B. The assessor performs a hands-on assessment of the auditees cloud services.
   C. This not a technical hands-on activity, although it can be informed by technical assessments.
   D. This exercise is conducted around a table—whether real or virtual—driven by a facilitator. Even if using tablets to facilitate the exercise, it is not hands on in the technical practitioner sense.

9. **A. Correct, also known as a first-party audit**
   B. Incorrect, this is a type of assessment, not an audit type
   C. Correct, also known as a second-party audit
   D. Correct, may also be certification or an attestation audit

10. Match the requirement to the correct requirements scheme.

Requirement	Requirements Scheme
A. Define SLA requirements	**2. Operational requirements** **Service level agreement is part of the operational requirements.**
B. Understand the applicable law	**3. Legal and regulatory requirements** **When setting legal requirements, part of the process is understanding and defining the applicable law.**
C. Define controls framework	**1. Security requirements** **As part of the security requirements process, the customer defines a list of controls—either generic (usually a third-party controls check list, i.e., CCM) or a custom list made by the customer.**

11. Match the terms and their definitions.

Term	Definition
1. Termination for cause	**D. Used to identify a potential consequence of bad behavior on the part of one party.**
2. Termination for convenience	**B. Used when circumstances have changed, and a party is able to stop without having to pay damages or penalty.**
3. Suspension	**A. Used when circumstances have changed, and a party needs time to reevaluate its options.**
4. Cure period	**E. Used when a party who has (or is accused of having) violated the contract is given an opportunity to correct the bad act.**
5. Extension	**C. Used to identify a contract that continues for a longer duration than originally expected.**

Chapter 2:

Cloud Compliance Program

Cloud Compliance Program

Cloud Compliance Program

2.1 Learning Objectives

After completing this chapter, learners will be able to:

1. Explain the fundamental criteria for cloud compliance programs.
2. Build and design a cloud compliance program.
3. Describe legal and regulatory requirements and standards and security frameworks.
4. Define controls and identify technical and process controls.
5. Recall CSA certification, attestation and validation.

2.2 Overview

This chapter covers specific aspects of a cloud compliance program, mainly from the perspective of a cloud service customer (CSC). The cloud service provider (CSP) perspective is also covered to clarify what a CSC should expect from a CSP in terms of compliance support and inheritance.

This chapter includes the following topics:

- Fundamental criteria for designing a cloud compliance program
- Guidance on how to build a cloud compliance program
- High-level framework
- Legal and regulatory requirements and standards
- Organizational requirements
- Control objectives, and control and process standards
- Technical and process controls
- Cloud-specific security control frameworks
- Cloud security certification and attestation

2.3 Fundamental Criteria for Cloud Compliance Programs

This section presents the fundamental criteria for a cloud compliance program.

2.3.1 CSP and Organizational System Considerations

The objective of a cloud compliance program is to ensure that all aspects of cloud adoption in an organization adhere to the requirements for which it is accountable. Enterprise risk management, governance policies, information security and data privacy are among the key components that drive decisions concerning the measures that the organization puts in place. Compliance ultimately aims to validate that these measures are appropriately designed, implemented and operated as intended.

Designing a cloud compliance program requires a complete understanding and documentation of the cloud ecosystem of the organization. Some cloud-ecosystem elements include current inventory of all cloud services, CSPs of services and delivery, and a responsibility matrix for all elements. The organization needs to determine its risk appetite prior to making any agreements with CSPs, because outsourcing of responsibilities to a provider may accompany an increased degree of shared responsibilities, including potential new risk to which the organization was not exposed before outsourcing. Maintaining an effective enterprise risk management program that syncs with the cloud service portfolio helps to prevent any surprises.

Other elements to determine during the design stage include establishing the systems that are critical to the organization and to the product or service it provides. Gaining this business understanding can involve several process-capture sessions to examine how the business units currently operate. It is important to document and maintain this information to ensure that the organization is aware of any gaps between the inherent and residual risk that exceed the organization risk appetite. These gaps may require special monitoring until fully remediated or accepted.

The nature and scope of the organization determine what is to be evaluated. This is especially important for organizations with business units that operate in various industries. Although all business units report to one parent or holding company, each industry can be subject to different regulations, laws or additional requirements that need to be met. Identifying and inventorying everything for which the organization is accountable strengthens the compliance program. Hence, as the organization compliance requirements increase in complexity, it is very important to align the governance model for the cloud compliance program to deal with the related complexities.

2.3.2 Geographic and Organizational Structure Considerations

From the cloud user's perspective, geographical considerations are also important (e.g., requirements for organizations in the United States are different from those for organizations in Europe). The nature of cloud services is that data often reside in multiple data centers in different locations for redundancy purposes. There are many laws affecting data storage and processing. These laws may require the cloud user to apply specific measures or controls to keep data secure for data privacy purposes. The choice of CSP determines the jurisdiction where data are stored. Depending on how the CSP organizes its infrastructure, the laws of several countries may apply. These considerations and variations of risk and compliance need to be built into the enterprise risk program. Specifically, relevant sovereignty risk needs to be included in the contract terms or conditions and in service level agreements (SLAs).

Another consideration is the audience of key stakeholders for the compliance program. The organization should answer the following questions:

- Who are the project sponsors, and what are their internal requirements for a cloud solution?
- What is the organizational structure, and how will communications be shared?
- Is there a command-and-control structure that may call for a specific set of controls and reporting?
- If reports are to be more widely distributed, does there need to be a different set of controls and different moderating?

There may be an independent division of the organization or a system that requires a specific level of compliance. By having strong documentation of all the environments that are under the scope of the compliance program, there is less room for error. Such documentation can include architecture diagrams, network flows, escalation processes, infrastructure builds, and role and responsibility models. The reporting model for the cloud compliance program needs to be congruent with the magnitude and complexity of the investment of the organization in cloud computing and the complexity of the entire ecosystem.

The different cloud models may have different compliance program configurations. Understanding the boundaries of the systems and their position within the organization can aid in how to validate or test that the controls in place are working properly. Aspects of the compliance program may need to be adjusted because of gaps being identified.

Under the shared responsibility model, which is an essential building block of contracts with a CSP, a distinction exists between security controls for which the provider is responsible and controls that remain the responsibility of the organization consuming the cloud service.

The cloud customer inherits certain controls from the CSP directly or through a third-party organization, such as a colocation vendor or IT services provider. Knowing what those controls are is essential to set up a compliance program (and later, for auditing it). Fundamentally, a compliance program determines who is responsible for what controls. Developing the methodology to test and validate the effectiveness of the operation and design of the controls put in place is critical, so the organization can rely on this methodology to monitor and report any drift or changes.

The fundamental criteria of building a cloud compliance program are summarized in the mind map in **figure 2.1**, and its components are further addressed in the next sections.

2.4 How to Design a Cloud Compliance Program

When designing a cloud compliance program, the organization should identify the key actors and requirements (**figure 2.1**), taking into account the following perspectives:

- Business/organization
- Cloud ecosystem
- Governance
- Risk

Figure 2.1—Mind Map for Designing a Cloud Compliance Program

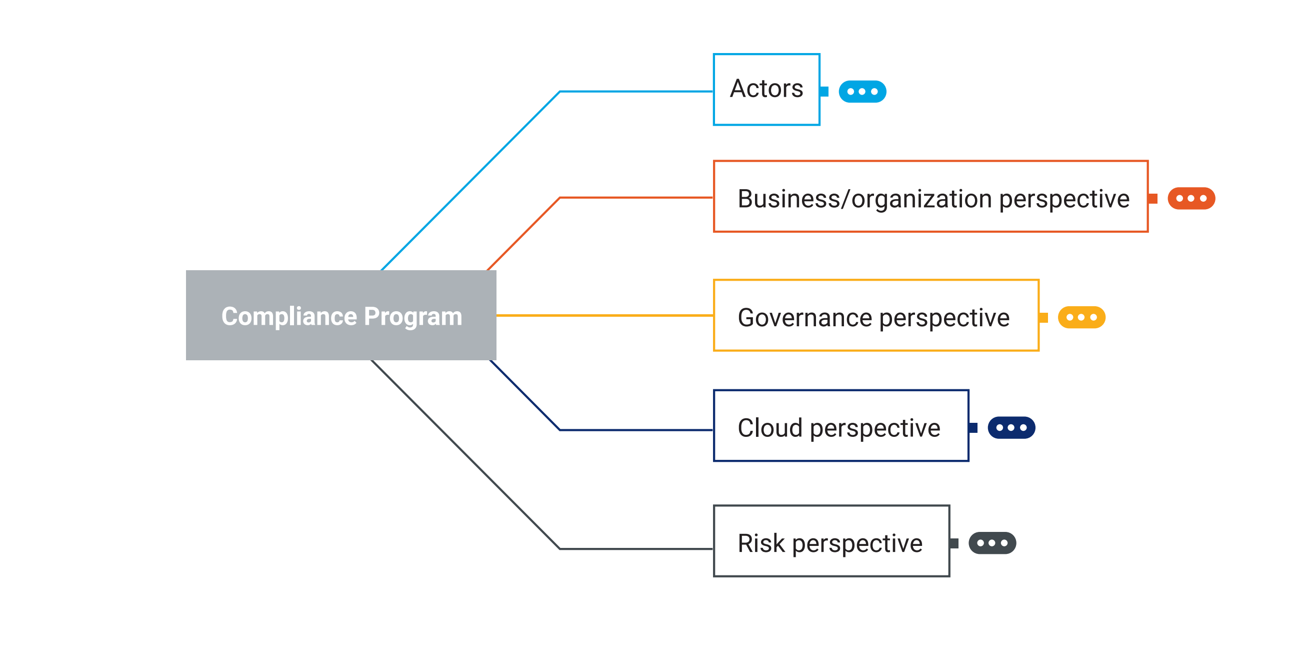

2.4.1 Key Actors

It is important to identify key stakeholders early in the design phase of the cloud compliance program. Stakeholders should include relevant decision makers, risk owners and executives accountable for business processes or objectives that are cloud dependent. Key stakeholders should also include anyone responsible for assessing, measuring or reporting on cloud compliance program performance. The key actors are the individuals who determine how the organization approaches the cloud; decide the strategy, budget, timeline, priorities, requirements and risk appetite; and decide how the technical implementation will be done.

It is important to interview and consult the key actors during the compliance program design phase (**figure 2.2**). Following are some of the most relevant questions to answer:

- Who owns the cloud strategy?
- Who owns the cloud policy?
- Who are the cloud service sponsors?
- Who are the cloud service portfolio owners?
- Who are the cloud service owners?
- Who are the data owners?
- Who manages the cloud service?
- Who are the control owners?
- Who are the service risk owners and managers?
- Who are the portfolio risk owners and managers?
- Who is responsible for service evaluation and approval?
- Who is in charge of compliance?
  - Who is the program owner?
  - Who builds the compliance rules?
  - Who checks and monitors the compliance rules?
- Are there any cloud brokers?
- What does the RACI (responsible, accountable, consulted and informed) matrix for each cloud service look like?
  - Who is responsible for what?
  - Who is accountable for what?
  - Who is consulted for what?
  - Who is informed about what?

Other individuals who should be consulted include, but are not limited to, internal auditors, information security, chief information officer (CIO)/chief technology officer (CTO), human resources (HR) personnel, external auditors, chief privacy officers and legal teams. Identifying the individuals who can answer these questions drives the development of the compliance program.

Figure 2.2—Mind Map for Key Actors in a Cloud Compliance Program

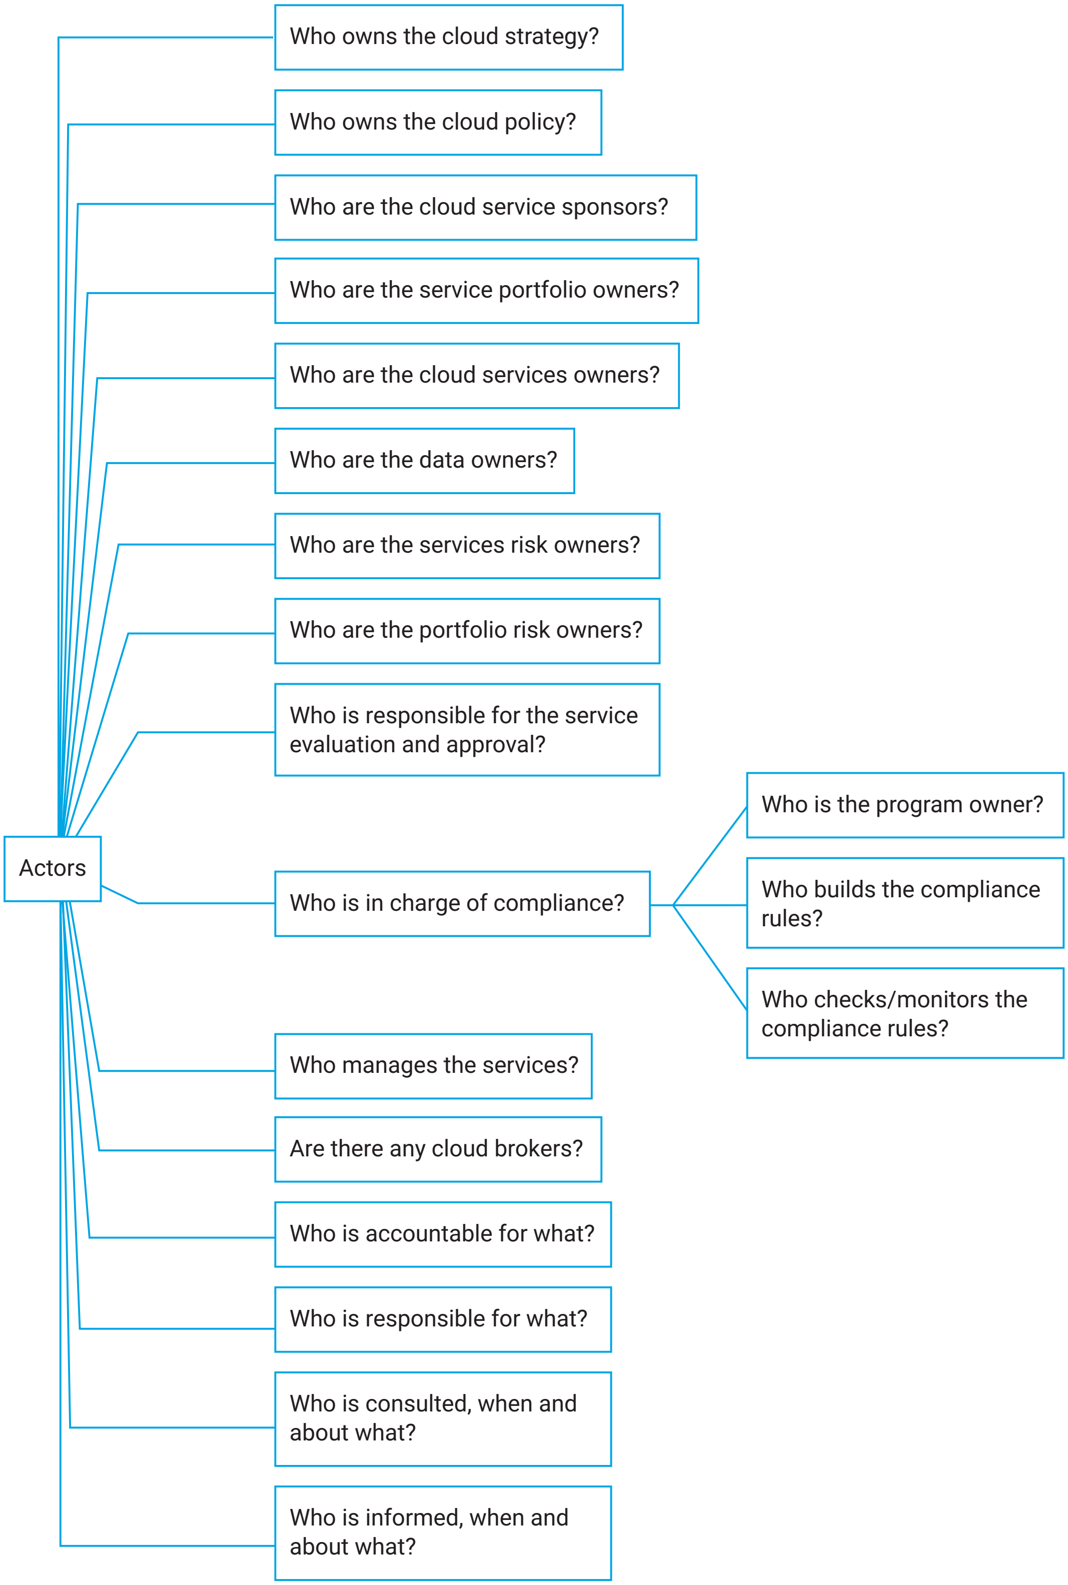
Who owns the cloud strategy?
Who owns the cloud policy?
Who are the cloud service sponsors?
Who are the service portfolio owners?
Who are the cloud services owners?
Who are the data owners?
Who are the services risk owners?
Who are the portfolio risk owners?
Who is responsible for the service evaluation and approval?
Actors
Who is in charge of compliance?
Who is the program owner?
Who builds the compliance rules?
Who checks/monitors the compliance rules?
Who manages the services?
Are there any cloud brokers?
Who is accountable for what?
Who is responsible for what?
Who is consulted, when and about what?
Who is informed, when and about what?

2.4.2 Business and Organizational Perspective

The cloud compliance program needs to consider different points of view, balancing the perspective of the business with the cloud computing perspective and the governance and risk management perspectives.

The business structure, strategy and approach determine which actions are required in the compliance program, because the type of organization determines what is strategically and tactically relevant, what resources are available, and what constraints, limitations and dependencies the business faces. In essence, the business structure, strategy and approach guide decisions on what is allowed and what is desirable.

In addition, the business and organizational requirements have a substantial impact on the governance, risk management and cloud computing strategy and approach (**figure 2.3**).

Figure 2.3—Mind Map for the Different Perspectives in a Cloud Compliance Program

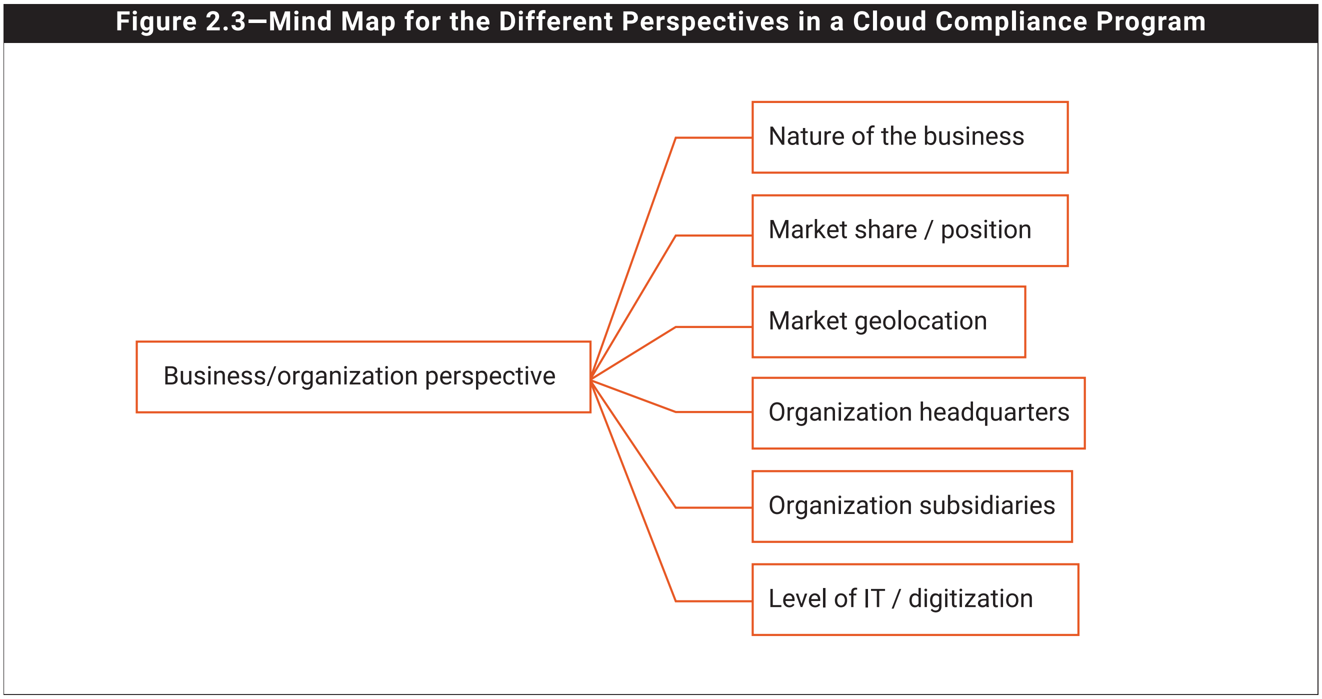
Business/organization perspective
Nature of the business
Market share / position
Market geolocation
Organization headquarters
Organization subsidiaries
Level of IT / digitization

Following are key questions to consider:

- **What is the nature of the business?**—In which business sector does the organization operate? Is it a highly regulated sector (e.g., banking and finance, or healthcare)? Is the organization or the sector considered part of the national critical infrastructure plan? Is the organization operating in a nonregulated sector?
- **Where is the company headquarters located?**—The answer to this question is one the most important variables to determine the applicable legal frameworks.
- **Where are the company subsidiaries?**—Does the company have other branch offices or subsidiaries? Where are they located?
- **What is the market geolocation?**—What markets are served by the company?
- **What is the level of IT dependency?**—Is the company heavily dependent on information and communication technologies? Can the company operate in the case of failure of the IT services? What is the level of digitalization?

Organization Status Quo and Preparedness

Understanding the past of the organization helps the team or individual in charge of cloud compliance to realize the history behind the current state of the organization compliance. Some key points to consider are:

- What is the current state of internal expertise in terms of audit and preparedness?
- What does the existing compliance program look like?
- Has the company gone through any audit in the past? If so, are the lessons learned applied today? (Is there a dedicated security program, internal or external, managing security decisions?)
- If the organization is multinational, does the compliance program make accommodations for independent country requirements? How does the organization ensure that it maintains compliance with independent country requirements?
- Has the organization been charged or penalized for previous compliance or legal violations?

Larger financial commitments that an organization makes in cloud computing can drastically impact the IT department. Many of the impacts have a direct effect on the IT risk profile. Ways that cloud adoption can potentially impact the IT department include establishing the need to take the following steps:

- Redesign the organization (possible reduction of staff)
- Reassess data governance requirements
- Amend relevant IT, cloud, security and privacy policies/standards
- Define skill and competency gaps
- Reassess tooling and product effectiveness
- Drive the use of cloud-native security tools and third-party management efforts
- Implement continuous monitoring
- Validate existing technical controls for their applicability

2.4.3 Governance Perspective

Cloud governance and cloud compliance are closely related. The mind map in **figure 2.4** provides key considerations from the governance perspective.

Figure 2.4—Mind Map of the Governance Perspective

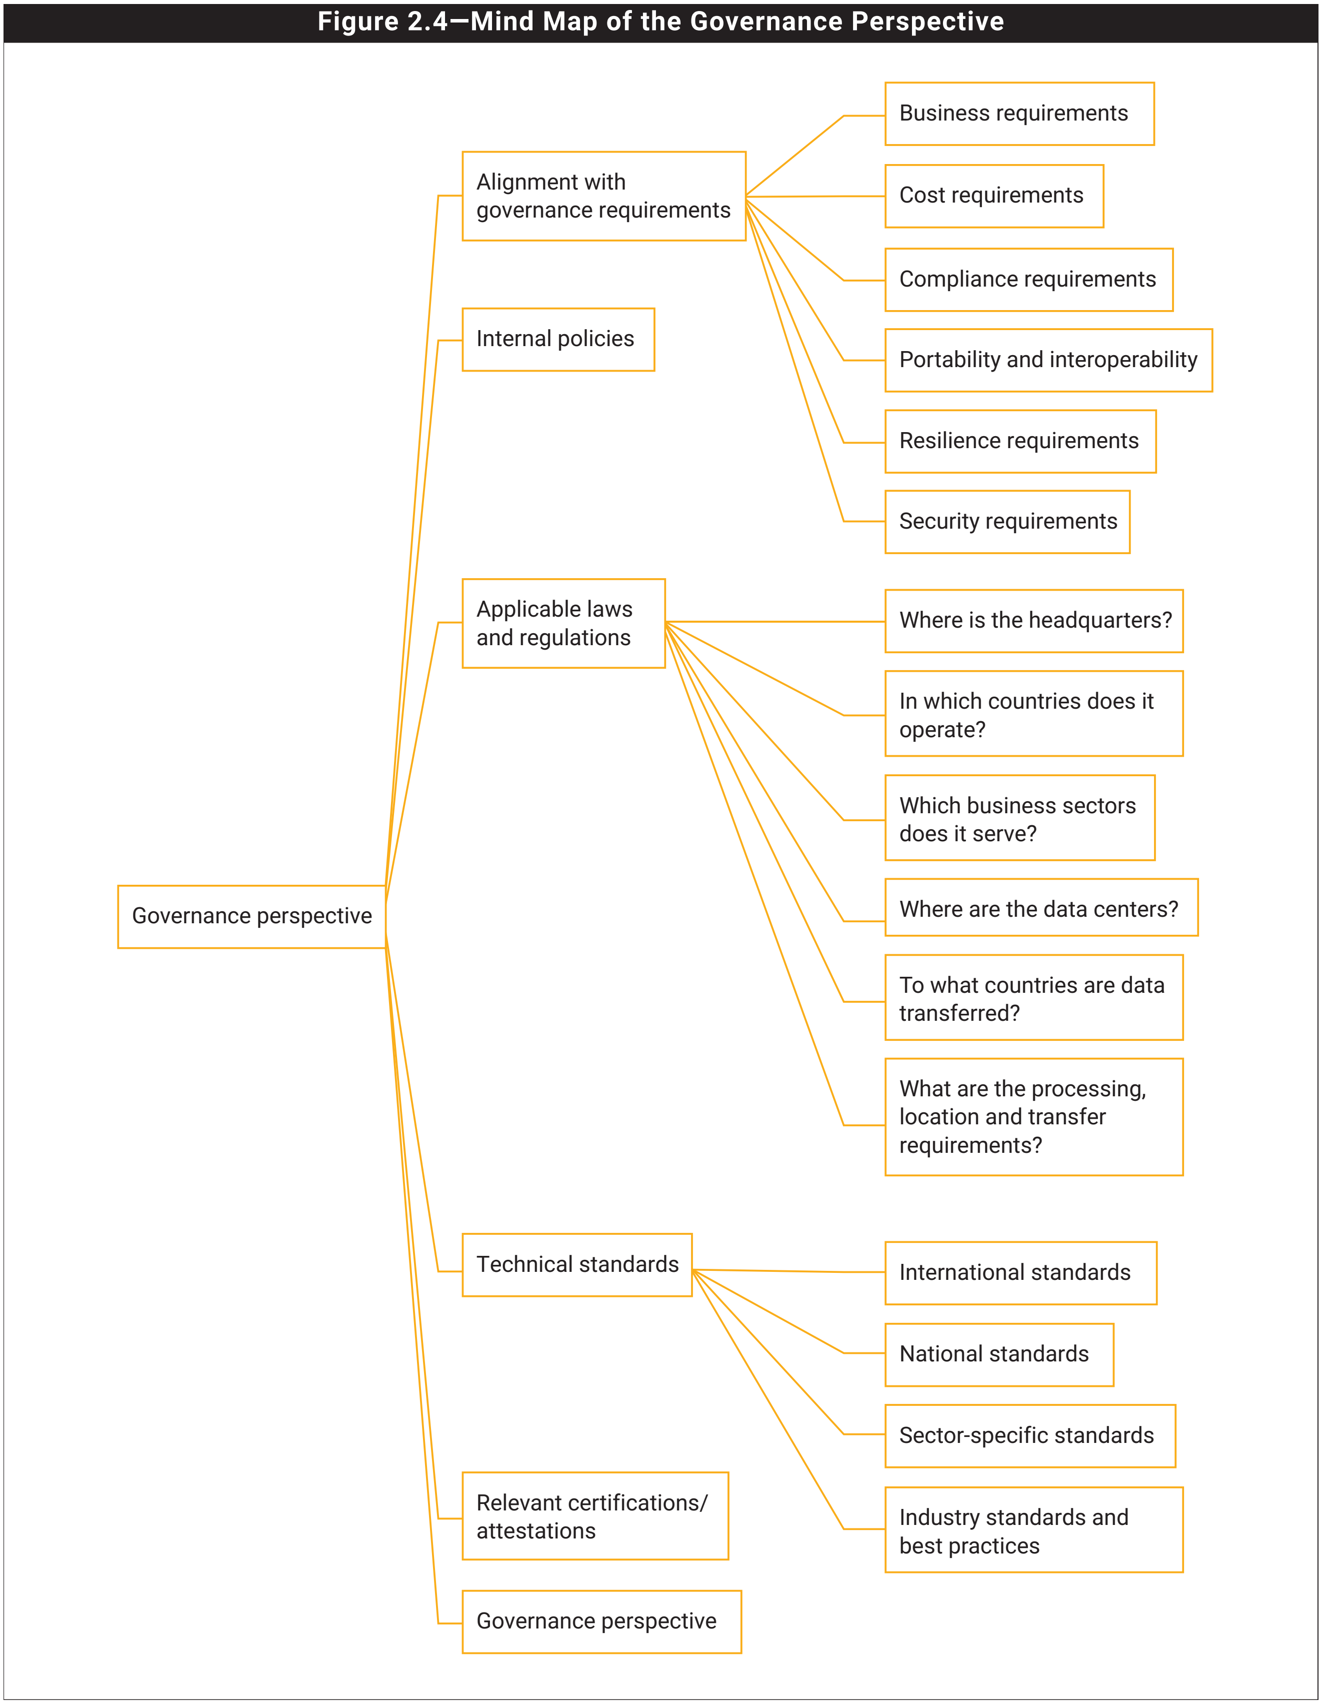
Governance perspective
Alignment with governance requirements
Business requirements
Cost requirements
Compliance requirements
Portability and interoperability
Resilience requirements
Security requirements
Internal policies
Applicable laws and regulations
Where is the headquarters?
In which countries does it operate?
Which business sectors does it serve?
Where are the data centers?
To what countries are data transferred?
What are the processing, location and transfer requirements?
Technical standards
International standards
National standards
Sector-specific standards
Industry standards and best practices
Relevant certifications/ attestations
Governance perspective

The team in charge of creating the cloud compliance program should consider the following:

- Alignment between the cloud requirements and the governance requirements:
  - What are the business requirements and how do they map to the cloud strategy and approach?
  - What are the budget constraints and the resources available?
  - What compliance constraints and priorities has the board imposed?
  - What are the security requirements?
  - What are the resilience requirements?
  - What are the privacy requirements?
  - What are the interoperability and portability requirements?
  - What are the exit strategies for key CSPs?
- Financial reporting:
  - Annual financial statement disclosure regarding IT and cloud risk
  - Key business and environmental risk
- Internal policies:
  - What are the policies for IT?
  - What are the policies for information security and privacy?
  - Does the organization have a dedicated cloud computing policy?
  - What does the policy say, and how does it align with the security and privacy requirements?
- Applicable laws and regulations:
  - Where is the company headquartered?
  - In which countries does it operate?
  - Which countries are served by the business of the organization?
  - Which business sectors are served?

Other key questions to ask include the following:

- Where are the data centers located?
- What types and categories of data are collected, processed and stored?
- In which countries are data transferred?
- What are the data processing, residency, retention and transfer requirements?
- How does the selected CSP/cloud service impact compliance with applicable laws and regulations?
- Applicable or relevant technical standards for security, privacy, interoperability and portability:
  - What are the relevant international standards?
  - What are the relevant national standards?
  - What are the relevant sector-specific standards?
  - What are the relevant industry best practices?
- Applicable or relevant international, national and sector-specific certifications and attestations:
  - What are the relevant certifications for interoperability and portability?
  - What are the relevant certifications for security, privacy, interoperability and portability?
  - What are the relevant attestations for security, privacy, interoperability and portability?

2.4.4 Cloud Perspective

Understanding the cloud computing context is a core challenge when designing and building a cloud compliance program, because it requires analyzing and understanding the approach of the organization from an IT and business perspective. Understanding the cloud computing context also means understanding the differences between legacy on-premises IT and cloud technologies in terms of compliance needs. A strong interconnection and interdependence exists between the decisions that an organization makes from the overall IT governance and risk perspective, and those from a cloud perspective.

For example, the organization risk appetite and budget constraints have a substantial impact on how aggressive or conservative its cloud approach is, which type of cloud services it uses, and for which type of data and business applications. Similarly, cloud strategy determined by the organization impacts the existing governance approach and the way it manages risk.

Differences in Legacy and Cloud Compliance Programs

The traditional risk management approach applies to cloud computing, but the tactics and implementation styles used for traditional IT versus cloud-based IT can be significantly different. Legacy compliance programs rely on systems with defined boundaries. With the introduction of cloud technologies, a major difference is the reliance on third parties for the delivery of technology resources, limiting the direct control and influence of the customer over the quality, availability and reliability of the services that a third party provides. Moreover, the same physical limitations exist (to the extent that the CSP attains the various compliance requirements), and at best are supplemented by virtual abstractions represented by software and automation. Hence, the concept of the cloud compliance program becomes the only vehicle available to validate the performance of the many cloud providers that may be fulfilling the compliance and business needs of the organization .

In a cloud or hybrid environment, the access network is public, the infrastructure is virtual, the perimeter is unknown and access can be gained from anywhere. The methods that an organization uses to access systems and data differ fundamentally in the cloud compared to legacy on-premises systems. The cloud, by nature, is a much more dynamic environment. Services are more likely to change, and they can scale up or down faster. Static compliance programs may not be sufficient to keep up with the rate of change with a cloud service. Therefore, the compliance program needs to align accordingly.

With a legacy program, the organization has complete control and responsibility for the infrastructure and all requirements that keep it secure. It is important to note that moving to the cloud does not shift all risk onto the CSP. Responsibilities are shared, based on the CSP model, but accountability remains with the customer.

Many organizations are now adopting applications that are multicloud in nature. These are applications built by one software as a service (SaaS) provider and hosted on the infrastructure of another provider. Designing for compliance in this cloud-to-cloud or multicloud environment is very different from designing compliance in a legacy environment. It introduces risk that is associated with supply chain visibility and validation.

Other considerations include how the rapid and dynamic rate of change in the cloud environment affects the risk posture of the cloud user. This consideration might raise questions about whether the cloud customer needs to change controls. The organization should have ways to monitor for change in the CSP service and drive adjustments if necessary. Due to the nature of cloud services that are easily consumed from a menu of services, such monitoring requires appropriate accommodation in the operation of the risk management program, so that it becomes much more sophisticated and agile as new services and features are adopted or modified using control panels and portals with minimal review or advance oversight. See **figure 2.5** for a mind map of the cloud perspective.

Figure 2.5—Mind Map of the Cloud Perspective

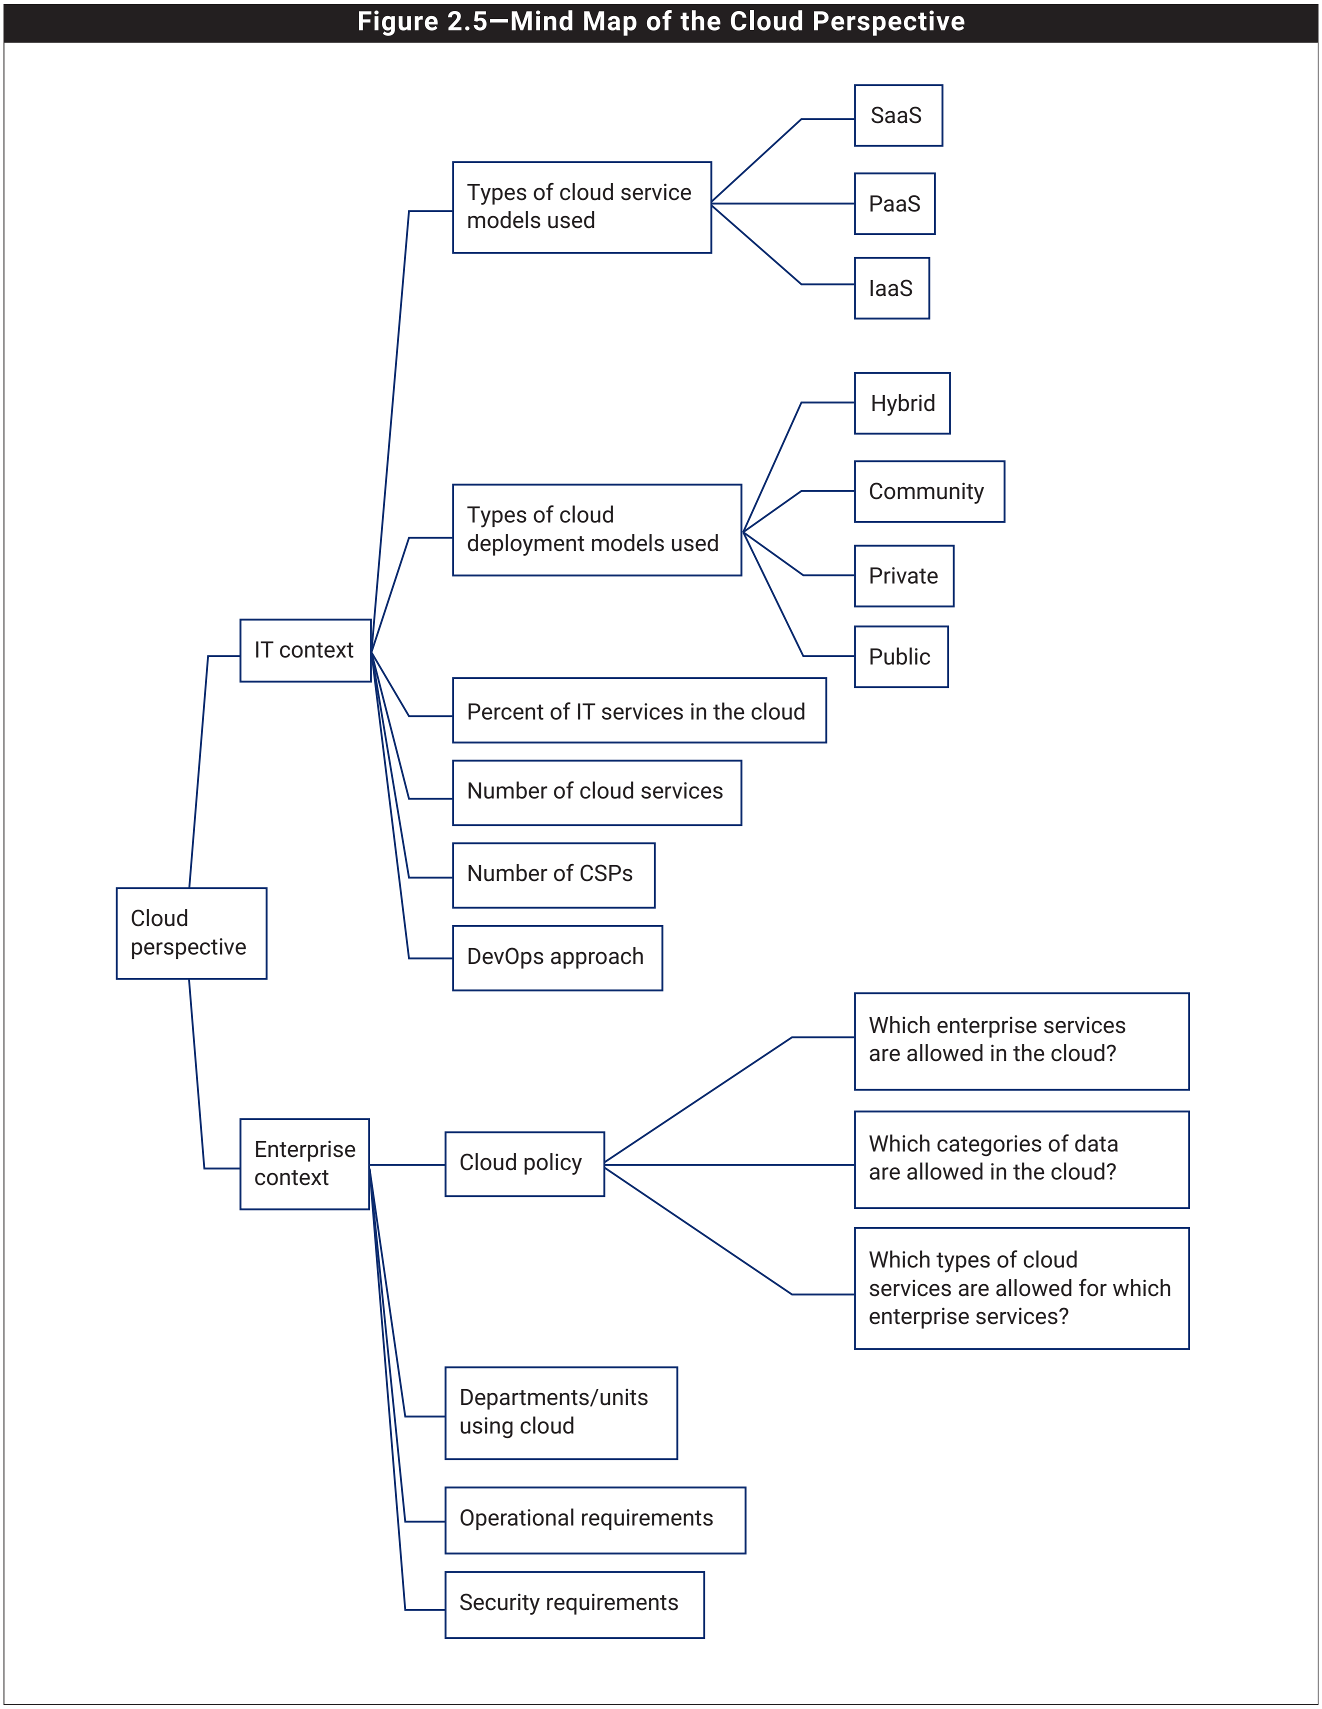
Cloud perspective
IT context
Types of cloud service models used
SaaS
PaaS
IaaS
Types of cloud deployment models used
Hybrid
Community
Private
Public
Percent of IT services in the cloud
Number of cloud services
Number of CSPs
DevOps approach
Enterprise context
Cloud policy
Which enterprise services are allowed in the cloud?
Which categories of data are allowed in the cloud?
Which types of cloud services are allowed for which enterprise services?
Departments/units using cloud
Operational requirements
Security requirements

Business Context of Cloud Approach

The cloud business context defines the scope, objectives, goals, policies and strategies of cloud adoption within an organization. Depending on its culture, the highest level of governance within the organization may define the cloud approach, e.g., with board and executive approval, or cloud adoption may be introduced with minimal or no approval or knowledge by the appropriate governance function (i.e., shadow IT). In the first case, with board knowledge, the organization should have a clearly defined cloud policy, strategy and road map. In the second case, the organization is likely to have an ad hoc opportunistic and tactical approach that lacks a defined cloud policy and strategy (clearly problematic from a risk management and compliance perspective).

Establishing a comprehensive cloud policy is very important to organizations that are planning to make major investments in cloud technologies and platforms. When creating policies, they should take the following considerations into account:

- Criteria for limiting services that may be moved to the cloud (e.g., only noncritical services or services that require high availability and resilience)
- Criteria for defining which data may be moved to the cloud and under what circumstances (e.g., nonsensitive and confidential data, only public data, or all data)
- Information security and privacy requirements specific to cloud services and deployments that are critical to defining the new business operational models
- Which deployment and delivery model is allowed for which type of data and services (e.g., critical business services and HR applications can only be moved into a private cloud; or public cloud environments can only be used for nonregulated data)
- How to define roles and responsibilities for technical and nontechnical personnel, including employees, contractors and business partners that may be using embedded cloud platforms and technologies
- Acquisition of the requisite software tools and services to manage and maintain the various cloud services and platforms, since on-premises solutions will not be viable
- Updates to procurement policies to ensure that proper risk assessments, product assessments and due diligence are applied to all cloud service acquisitions, including the creation of standard service level agreements and contract addendums or exhibits
- Empowerment of the cloud compliance program to review and approve all cloud services and cloud service providers
- Criteria for defining responsibilities for enforcement and potential consequences of noncompliance (an important consideration for developing any new policy)

It is also important to consider who is using the cloud and to know the business, operational and security requirements.

Different business units and departments may have different business and security requirements. Consequently, they may be using cloud services for different purposes or may not be using the cloud. This can create governance tension within the organization and, consequently, impact the cloud compliance program.

For example, the marketing team in a multinational company may be using a CSP that does not specify to which countries the data are transferred and stored. This can conflict with legal and privacy requirements if the organization is being held accountable for adhering to geographic regulations, such as the EU General Data Protection Regulation (GDPR).

IT Context

The IT infrastructure is the core variable to consider when designing a compliance program. The shape that the program takes depends on its configuration. The basics are defined by the mix of service deployment models (private, public, hybrid and community) and service delivery models (IaaS, PaaS and SaaS).

The typical scenario is a hybrid cloud model: The cloud customer integrates the public cloud services into its existing on-premises infrastructure or private cloud. From the delivery perspective, it is likely that the customer has one or two IaaS or PaaS providers and multiple (from ten to thousands) SaaS providers.

It is important to consider who is managing the IaaS and PaaS. Is it the cloud customer itself, or does it outsource this role to a managed CSP? If it is outsourced, then managing the subcontractors becomes a key consideration in the compliance program.

The compliance program is also influenced by the level of maturity and pervasiveness of the cloud implementation, and whether the organization is adopting a cloud-native development approach.

Important questions to consider are:

- In which phase of the cloud journey is the organization? Is it developing a business case, introducing cloud for noncritical services or trying to understand the extent of the shadow IT issue, or does it already have several years of cloud experience?
- Is IT centralized within the organization?
- What is the ratio between the total number of IT services within an organization and the number of cloud services? Is the company mostly cloud-based or mostly on-premise?
- How many CSPs are used? What underpinning contracts are in place?
- What is the total number of cloud services used? Tens? Hundreds? Thousands?
- How is cloud-service use monitored? Are there mechanisms in place to understand if a service that originally was contracted for one purpose is evolving over time? In general, it is paramount to monitor the cloud service during its entire life cycle to ensure that compliance requirements are enforced during the initial service evaluation, during the execution and during its retirement (particular attention should be given to what happens to the data after a service is terminated).

Two variables that have dramatically changed the face of compliance in the last few years are DevOps and automation. For organizations that are embracing cloud-native development processes, the approach to auditing and compliance is rapidly evolving toward a more agile way to implement and enforce compliance requirements. The pace at which new services are delivered, features are added, and changes take place necessitates a new way to structure compliance. A new way also offers the possibility for a more effective approach to compliance based on automation. Compliance becomes embedded in the development pipeline according to the idea of compliance-as-code. See chapter 8 for in-depth treatment of the topic of DevOps and continuous compliance.

2.4.5 Risk Perspective: Risk Appetite and Risk Assessment

Identifying and understanding the risk to an organization is the foundation of designing an effective compliance program. Section 1.4.10 addresses risk management, defining some cloud-specific risk and providing an overview of key components of a risk management approach, based on best practices from the European Union Cybersecurity Agency (ENISA). The focus of this section is to make the connection between the risk management approach and the cloud compliance program to understand how they influence each other (**figure 2.6**).

Figure 2.6—Mind Map for Risk Appetite and Risk Assessment

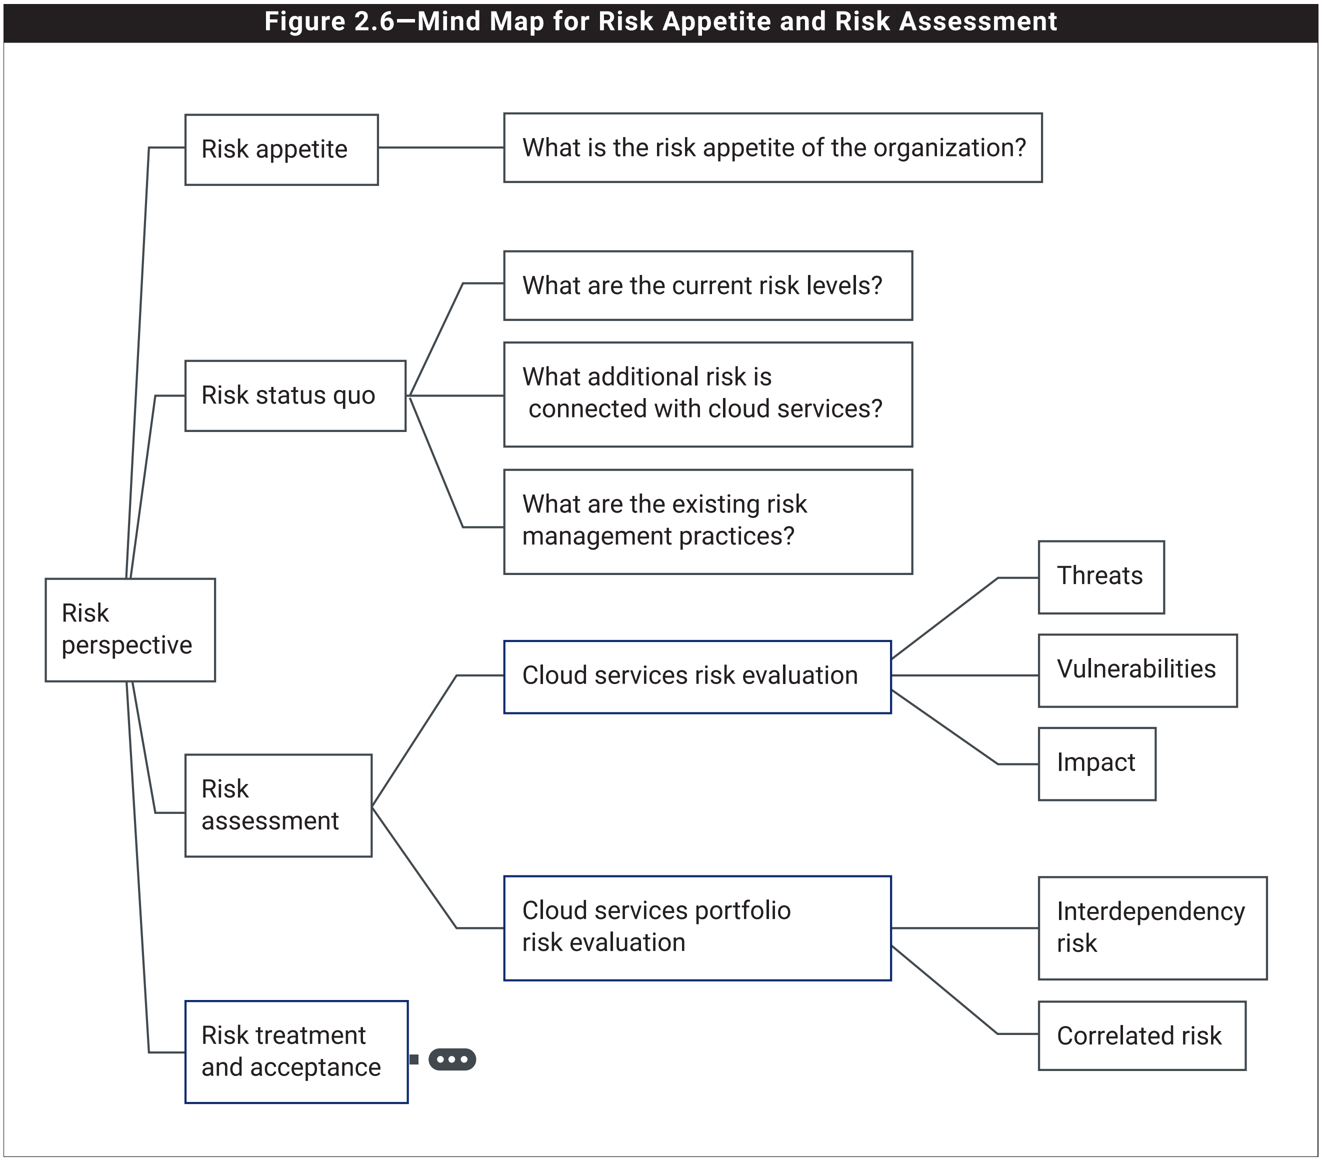
Risk perspective
Risk appetite
What is the risk appetite of the organization?
Risk status quo
What are the current risk levels?
What additional risk is connected with cloud services?
What are the existing risk management practices?
Risk assessment
Cloud services risk evaluation
Threats
Vulnerabilities
Impact
Cloud services portfolio risk evaluation
Interdependency risk
Correlated risk
Risk treatment and acceptance

Risk Appetite and Status Quo

The risk appetite of an organization determines its risk management style. An organization that is a risk taker may be willing to accept more risk in the face of a higher return (e.g., a shorter time to market for a certain service, or adoption of IaaS providers that do not provide the required security capabilities, or adoption of a SaaS service without a thorough security assessment). In contrast, a risk-averse organization minimizes its risk exposure to lower the probability of suffering a negative event, but at the cost of losing possible opportunities. A difference in the terms of risk appetite might also appear between sections, departments and business lines of the same organization. For example, within financial services organizations, which are typically risk-averse, there are business guidelines for which time-to-market initiatives take priority over information security. The design of the compliance program reflects the trade-offs between risk and opportunities, which themselves determine the compliance risk that the organization is willing to take.

Although the risk appetite defines the modus operandi, the organization status quo defines the position from which the cloud compliance program starts. Analyzing the status quo helps to determine the existing level of risk, the additional perceived and real risk connected to adopting a cloud service, and the existing mitigation strategies and practices.

In general, if an organization already has an established and mature IT risk management and compliance program, the cloud component is likely to be an evolution of that, with all the existing pros and cons attached. The existence of a mature IT compliance program offers the cloud program know-how, culture and capabilities on which to build, but it may create some friction between approaches to auditing that may be substantially different, especially if the cloud team embraces cloud-native development approaches like DevOps.

When adopting new technologies that have the potential to significantly impact business progress and competitiveness, the most important thing for the organization is to have the capability to fully comprehend the risk and potential outcomes associated with the technology decisions that it makes. Hence, the organization should introduce a cloud compliance program carefully and deliberately, with confidence that the business, IT and security functions have the competency, skill sets and authority to effectively guide its cloud journey.

Risk Assessment

Risk assessment methodologies in the world of information security work broadly, as follows:

1. Identify the scope of the information system and its governance.
2. Identify assets (e.g., intellectual property, customers, websites).
3. Identify threats to those assets, and threat agents (e.g., hackers, disgruntled employees, competitors).
4. Identify vulnerabilities that threat agents may exploit to realize a threat.
5. Compute risk as the product of the likelihood of the realization of threats and the impact of their realization.
6. Evaluate the computed risk against the risk appetite so that the organization takes steps based on the acceptability criteria.

The intent of completing a risk assessment is to determine potential threats and vulnerabilities, and their likelihood, and their impact should they occur. The output of this process helps to identify appropriate controls for reducing or eliminating risk, a topic in the next section. The results of the risk assessment process are a key component of designing the compliance program, because they inform the risk registry.

2.4.6 Risk Perspective: Risk Treatment and Acceptance

This section discusses risk treatment and acceptance approaches and how they influence the design of a compliance program. After an organization identifies and assesses risk, the organization needs to determine the measures to put in place to achieve an acceptable risk level (see **figure 2.7**). That means identifying security and privacy control objectives, designing technical countermeasures (technical controls) tailored to the organization, and implementing and monitoring their performance. All of that is typically done with the support of tools. From the business standpoint, it is fundamental to achieve a clear understanding of due diligence responsibilities and liabilities. In the context of cybersecurity, due diligence refers to the process of identifying and remediating risks associated with third-party vendors.

Figure 2.7—Mind Map of the Risk Perspective

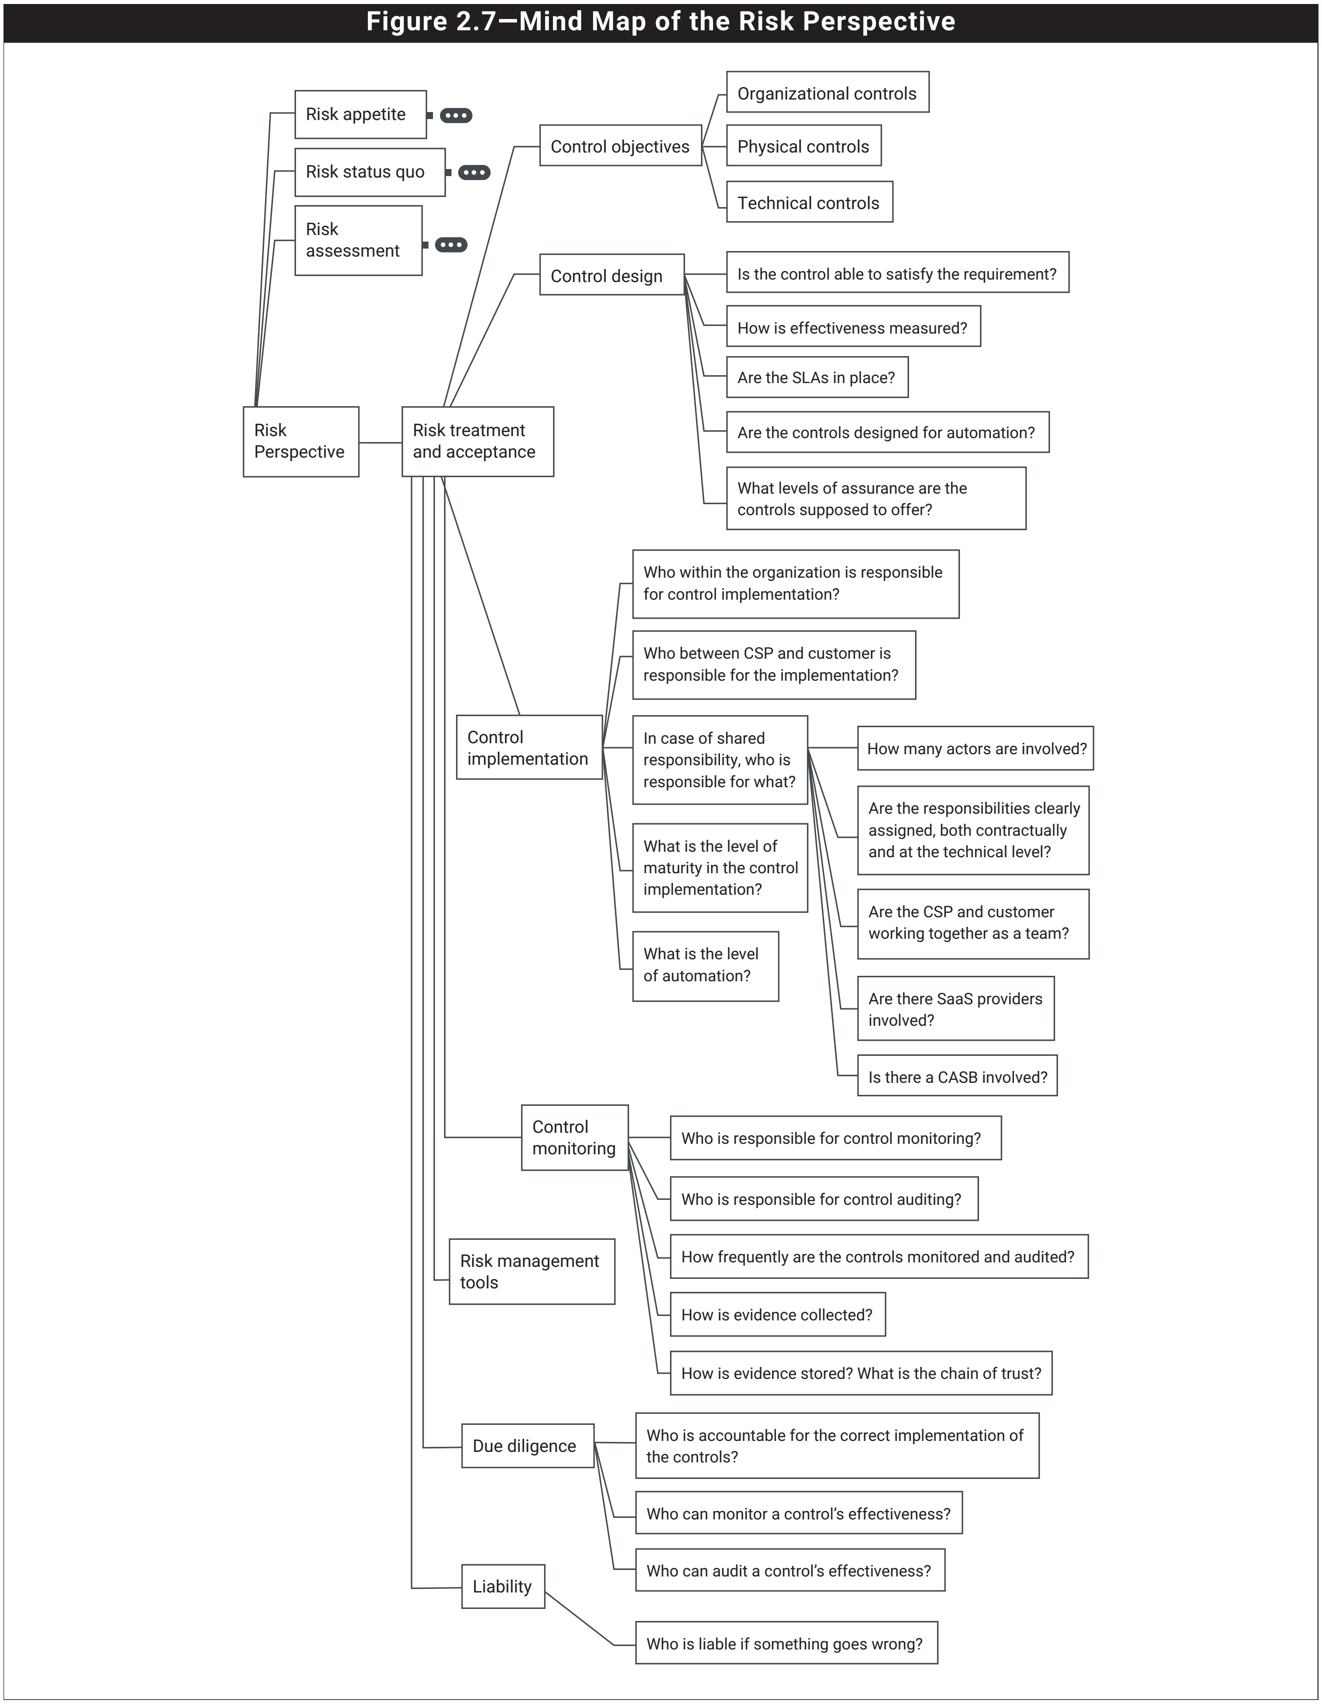
Risk Perspective
Risk appetite
Risk status quo
Risk assessment
Risk treatment and acceptance
Control objectives
Organizational controls
Physical controls
Technical controls
Control design
Is the control able to satisfy the requirement?
How is effectiveness measured?
Are the SLAs in place?
Are the controls designed for automation?
What levels of assurance are the controls supposed to offer?
Control implementation
Who within the organization is responsible for control implementation?
Who between CSP and customer is responsible for the implementation?
In case of shared responsibility, who is responsible for what?
How many actors are involved?
Are the responsibilities clearly assigned, both contractually and at the technical level?
Are the CSP and customer working together as a team?
Are there SaaS providers involved?
Is there a CASB involved?
What is the level of maturity in the control implementation?
What is the level of automation?
Control monitoring
Who is responsible for control monitoring?
Who is responsible for control auditing?
How frequently are the controls monitored and audited?
How is evidence collected?
How is evidence stored? What is the chain of trust?
Risk management tools
Due diligence
Who is accountable for the correct implementation of the controls?
Who can monitor a control's effectiveness?
Who can audit a control's effectiveness?
Liability
Who is liable if something goes wrong?

Control Objectives

There is an argument that compliance in the cloud environment is theoretically the same as compliance in the legacy IT context. They carry similar types of risk, and many of the same preventive and detective controls are applicable. The key differences are the rate of change, the access from anywhere that the cloud allows, and some variations in terminology.

There are generally three types of controls: preventive, detective and corrective.

Examples of each type of control are:

- Data encryption (preventive)
- Monitoring (detective)
- Incident response (corrective)

These controls may need to adapt and evolve depending on changes in the threat landscape. The compliance team should answer these questions:

- Over a given period, how accurately does the compliance program reflect the latest threats to the business?
- Have threats led to any changes in the choice and deployment of controls?

Additional details on control objectives are presented in later sections.

Control Specifications Design

With cloud services, the need to design controls into processes and services at the outset is vital, because these services are primarily accessible via the Internet rather than behind a corporate network, and, therefore, the risk is greater. Control specifications design is the way that those business processes are designed. It considers the objective of the process, how the process achieves its result, and process inputs and outputs.

It is the job of the architect and the engineer to identify the technology components that will either embed or deliver particular aspects of the business process and determine how to orchestrate them. This means determining how they all work together to deliver the required functionality to the required standard (e.g., an approved list of technologies or vendors). This can come in the form of an architectural blueprint that defines how the architect and engineer put things together.

When designing processes and services, a group is normally responsible for making key decisions about how that process or service is designed, and for considering the trade-offs between cost, time, convenience, security and privacy. Several important questions to address include:

- Is the control able to satisfy the requirements?
- How is effectiveness measured?
- Does the organization define service level objectives (SLOs) and service qualitative objectives (SQOs)?
- What is the process for defining SLOs and SQOs?
- How do internal SQOs and SLOs map to the SLAs that the CSPs offer?
- Are the controls designed for manual implementation or for automation?
- How do cloud-native development approaches (i.e., DevOps) impact the control design process?
- What is the expected level of assurance that the control is supposed to provide?

See the "Defining Controls" section in this chapter for additional details on controls.

Control Implementation

After controls are designed to align with the organizational architecture and can satisfy the risk requirements, controls can be implemented. The most important factor to consider when implementing cloud-related controls is the shared responsibility model. Key questions to consider when designing the compliance program follow:

- Who is responsible—the customer or the CSP?
- Is there a responsibility assignment (RACI) matrix to determine who is responsible, who is accountable, who is contributing and who should be informed?
- If the CSP is responsible, who within the organization should determine if additional compensating controls are required? Who should implement them?

When considering shared responsibilities, the following questions help determine who does what:

- How many actors are involved?
- Are the responsibilities clearly assigned contractually and at the technical level?
- Are the CSP and the customer working together in a joint team?
- Are there SaaS providers involved?
- Are there cloud access security brokers (CASBs) involved?
- Are other Security as a Service (SecaaS) providers involved?

Lastly, the organization needs to determine the level of the maturity of the control implementation. It is helpful to answer the following questions:

- Are the controls implemented on an *ad hoc* basis, with a lift-and-shift approach,[31] and no automation in place?
- Are cloud-specific controls in place? Are there simple automation techniques in place?
- Is a cloud-native approach used across the organization?
- Are there guardrails in place?
- Is the system optimized and automation centralized?

Control Monitoring and Compliance Reporting

Risk management and compliance are not one-off exercises, but a process of continuous improvement that requires checking and monitoring of controls and reporting the results back to the key stakeholders. Some key questions to be addressed in the compliance program design phase follow:

- Who is responsible for the control monitoring?
- What extra activities does the organization undertake for cloud computing compared to on-premise?
- What changed after the organization started using the cloud?
- What is the level of maturity of the monitoring process?
- Who is responsible for auditing?
- How frequently are the controls monitored and audited?
- How is evidence collected?
- How is evidence stored? What is the chain of trust? What are the differences between cloud and on-premise?

[31] The lift-and-shift approach to cloud migration entails that workloads, applications and data will migrate from existing on-premises infrastructure(s) to the cloud with minimal or no changes.

The organization must have some method for delineating compliance reporting. This method should capture what is being reported (e.g., which metrics are used), to whom it is being reported, specific timelines and an indication of what the organization is doing with the data it handles to improve its environment and reduce its risk. Example compliance reporting questions follow:

- How will the organization measure the success of its compliance?
- With whom (stakeholder and business units) is the data shared?
- What is the frequency of the reporting? (Applicable laws and regulations largely influence this variable.)
- What is the process for putting lessons learned into action?
- Is the organization responding quickly enough to an incident to minimize the risk?

Tools

Depending on its size, complexity, budget and maturity, an organization might choose to adopt tools to streamline and improve its risk management and compliance programs. With the cloud compliance and governance markets evolving quickly, a selection of solutions have become available to address and report status more quickly and efficiently. Governance risk and compliance (GRC) tooling, such as GRC tools or cloud access security broker (CASB) solutions take the manual efforts from managing spreadsheets and other compliance processes and automate them into workstreams of day-to-day tasks.

These solutions are either on premises or cloud based, and they enforce security policies between cloud service providers and consumers. It is essential to combine and link the requirements of the organization with the cloud-based resources. Having real-time, consistent tools that handle this risk and organizational enforcement may help ease the minds of those responsible for its security posture.

Due Diligence and Liability

The need for due diligence is crucial for any organization and particularly for one that relies heavily on possibly untrusted third parties. Cloud computing is a business model that blurs the boundaries between customer and provider and introduces a complex supply chain involving several untrusted partners. Therefore, organizations should give additional emphasis to the due diligence approach when designing a compliance program.

Following are fundamental questions to address:

- Does the organization have an accountability program?
- What does it look like?
- How is accountability for the correct implementation of the controls enforced?
- How is the internal auditing function organized?
- How is accountability for the service evaluation managed?
- How is accountability for the cloud service life cycle managed?
- How is liability attributed?

2.5 How to Build a Cloud Compliance Program

This section explains how to build a cloud compliance program.

2.5.1 Components of a Cloud Compliance Program

A cloud compliance program consists of a set of recurring activities that an organization undertakes to assess its implementation of an information system management program against a compliance regulation or framework (high-level framework). The cloud compliance program must consider the high-level framework, the perspective of the organization and its risk management objectives.

The compliance program includes the following:

- High-level framework
- Business and organizational perspective
- Types of audits authorized, and schedule or trigger events
- Enterprise policies and procedures
- Documented information security activity review
- Risk management and risk assessment process
- Auditing and assessment tools
- Any mandatory audit-related training, education and awareness
- Dedicated compliance team/personnel

When developing a cloud compliance program, the cloud customer must review the cloud services it deploys and determine how those services fit within its policies and procedures to prevent, detect, contain and correct security violations.

High-level Framework

When building a compliance program, what specific frameworks should the organization consider? The following factors to consider relate to the nature of the business activities of the organization:

- Organization location and that of its customers, which determine the applicable legal frameworks and industry standards
- Organization code of conduct and best practices
- Dynamics of the market of reference of the organization
- Tactical and strategic plans
- Existing compliance program and its level of maturity (or lack of one)
- Risk program

One way to start the process of defining the high-level framework is to identify the key questions to ask, the relevant structures and units involved, and the key business stakeholders who should provide answers and make decisions.

In other words, define the cloud compliance requirements and determine their interplay with the business strategy, goals and other compliance requirements of the organization. The effectiveness of the process of understanding and mapping the key individuals who are responsible, accountable and possibly liable within the decision-making process; who own the risk; and who possess the information and data that should drive the compliance goals determines the success or failure of the cloud compliance program.

After the organization identifies those key requirements and stakeholders, it may choose to build its compliance program using frameworks from established international information security standards, such as ISO/IEC 27001,[32] and the NIST Cybersecurity Framework.[33] See section 2.6 for additional details on relevant security standards.

Types of Audits

The organization must define the types of audits it will accommodate or support. Audit types include the following:

- **Governance and strategy audit**—This type of audit evaluates the framework of the organization for defining requirements, performing risk assessments, monitoring controls, reporting adherence and developing a strategic plan for the cloud.
- **Configuration and activity monitoring**—This includes logging, monitoring, scanning and alerting of a system, account or environment. It can be achieved using real-time automated scripts or manual testing. Organizations continuously perform this type of audit as part of operations, helping them to implement continuous assurance and continuous compliance as described in sections that follow.
- **Access review**—This is the review of all user and service accounts, and permissions within established information system boundaries, including on-premises systems, cloud environments and other applications. As a result of each review, unused or invalid access is disabled.
- **Compliance and controls audits**—This is the audit performed against technical, administrative and physical controls, as defined in the organization policies and procedures. Typically, either the designated internal audit team or external audit firm performs IT or cloud audits following a potential trigger event. These events might include:
    - Scheduled compliance audits
    - Business associates' requests or complaints
    - Discovery of significant vulnerabilities or breaches

Policies and Procedures

Policies are high-level statements that define the rules by which the organization staff must abide while they carry out their various responsibilities. Procedures provide step-by-step instructions for specific routine tasks and may even include a checklist or process steps to follow.

Policies and procedures are usually written to document the high-level direction of an organization. They reference the relevant standards to follow and the processes and procedures necessary for implementation and operation. Policies should include actions to take when violations are detected, including personnel-related actions and reporting.

Documented Information Security Activity Review

The information security activity review encompasses the review and analysis of security logs, events, and audit trails by the security team, with the assistance of automated systems and processes. The review should look for information security activities, such as the following:

- The system must be configured with correlation rules and policies to identify suspicious activity, vulnerabilities and misconfigurations.
- Rules applicable to alerts and escalations are defined and implemented to notify the responsible staff.

[32] ISO, ISO/IEC 27001:2013 *Information technology – Security techniques – Information security management systems – Requirements*, October 2013, www.iso.org/standard/54534.html

[33] National Institutes of Standards and Technology, "Cybersecurity Framework," www.nist.gov/cyberframework

- Incidents are documented and escalated to a risk management process and logged in the risk registry.

The Risk Register

Risk assessment and risk management activities are documented in the risk register. This is a document or software that includes all the risk and threats identified during configuration and activity monitoring, access reviews, and compliance and controls audits.

It is very important to build and maintain a risk register as a software platform or a document. It records all the risk captured by the security committee that has an impact on business processes, applications and security. The risk register records the outcome of quarterly vulnerability assessments and remediation plans, for example. The subcomponents of a risk assessment and risk management plan follow:

- Risk analysis plan
- Process to review and report audit findings
- Process to remediate controls deficiencies
- Risk register

Auditing and Assessment Tools

The audit and compliance team may select and use assessment tools to detect vulnerabilities and intrusions. These tools may include the following:

- Scanning tools and devices
- Password cracking utilities
- Network sniffers
- Security agents installed locally on servers and endpoints
- Vulnerability scanning software
- Passive and active intrusion detection systems
- Penetration testing tools.

Audit-Related Training, Education and Awareness

Some regulations require that the workforce receive specific training on a periodic basis (usually once a year). To comply with the mandate, an organization maintains a list of the regulation and the required training, and a record of workforce attendance and participation in the training program. Training exercises can help organizations prove that their employees were made aware of their responsibilities and applicable sanctions or corrective disciplinary actions, if the audit process detects a failure to comply with the requirements.

Dedicated Team/Personnel Assigned the Job Function/Responsibility of Security and Compliance

The organization must assign a dedicated team or personnel to the job function of security and compliance to enforce the principle of segregation of duties. Segregation of duties can be achieved via a combination of assignment of roles and responsibilities to different personnel, and automated enforcement for software-defined processes. Such segregation of duties and related review and approval processes ensure checks and balances. When applicable, designated personnel, separate from the individuals who are performing the work, must provide reviews and approvals.

2.5.2 How to Integrate Cloud Services Into the Company Compliance Program

When integrating all compliance aspects of an organization with its portfolio of cloud services, consider the answers to some important questions, such as the following:

- How many cloud providers serve the organization?
- What are the relevant measures of on-premises technology vs. cloud platforms use?
- Which entity within the organization is accountable for the performance of the CSP?
- Which regulator or legal entity sets the compliance requirements?
- Are the regulator and applicable laws and requirements foreign or domestic?
- Is the organizational compliance ownership centralized or decentralized?
- Does the organization maintain an enterprise risk management program (ERM)?

The integration model should conform to the governance model of the organization. This conformance becomes more critical with increases in the number of organizational entities and the number of cloud providers.

As illustrated in **figure 2.8**, the organization should prepare and maintain a RACI chart or its information equivalent to document the compliance responsibilities and ownership of accountability. Such information requires collaboration with relevant stakeholders and demonstrates organization-wide communication of compliance requirements and evidence of monitoring.

This methodology and approach can provide a holistic and seamless view of the enterprise responsibility for compliance with prevailing laws and regulations. This design and evidence of effective reporting and monitoring should be subject to review by the internal audit function and examined on a periodic basis, as determined by the enterprise risk management program. In addition, US-based organizations must adhere to the unique requirements of corporate compliance programs, evaluating the reporting and monitoring by design and operation, against the organizational guidelines of the US Sentencing Commission.[34]

A key consideration when evaluating the compliance program, specifically for CSPs, is the ability of the organization to identify and include all providers. The undocumented existence of shadow IT or data cloud providers can present a major unmeasured risk to the organization.

Figure 2.8—Matrix of Compliance Ownership vs. Cloud Service Provider Ownership

	GRC Department IT	Business Unit IT	Corp Info Security	Corp Compliance	Business Unit Compliance	Corp Privacy
CSP (1)	A	I	R	C	I	C
CSP (2)	I	A	R	C	C	I
CSP (3)	I	I	A/R	C	I	I
CSP (4)	I	I	C/I	C/I	R	A
CSP (n)	I	I	I	C/I	A/R	I
Legend: **A** = Accountable, **R** = Responsible, **C**= Consulted, **I** = Informed						

As a result of the concepts that cloud computing introduces, such as shadow IT, shared controls, cloud supply chain risk and CSP configuration features, traditional governance may no longer be as effective in achieving enterprise-wide compliance. This decrease in effectiveness is due to the accompanying business transformation that typically happens in tandem with the massive adoption of cloud platforms. As organizations rely on technology to disrupt current industry practices with innovative ways to conduct business, they frequently modify the manner in which they achieve or report legal and regulatory compliance. In addition, they accomplish compliance processes and

[34] United States Sentencing Commission, "Organizational Guidelines," www.ussc.gov/guidelines/organizational-guidelines

reporting through automated techniques that are less transparent and more difficult to validate. This is especially the case when the compliance actions and evidence of operation are in the cloud and not performed with on-premises solutions. This decrease in effectiveness in achieving enterprise-wide compliance has been shown repeatedly in the marketplace, with the inappropriate sharing of personal information within the CSP ecosystem, the limitations of relying on SaaS-provider attestation reports for coverage of compliance controls and when compliance conformance is the CSP responsibility.

This effect is partly attributable to the customer's lack of contracting power and partly to the inability of the enterprise compliance function to understand, detect and mitigate its loss of visibility.

Hence, the ability to ensure that organizations maintain accountability as they transfer their responsibility for compliance with various regulatory requirements to their CSPs may be challenging. This ability requires an understanding of the nature and scope of the risk, accompanied by enhancements in the organization governance model. Greater integration between the compliance oversight that the IT department provides and the various other business unit and corporate-level compliance functions is required.

2.6 Legal and Regulatory Requirements, Standards and Security Frameworks

Having a clear understanding of external requirements that should be contemplated in the cloud controls program is critical. Influential sources of external requirements are commonly based on key factors associated with the organization, such as industry sector, geographic location and size. Although, in some cases, these laws and regulations may be cloud-specific, they usually apply more broadly, regardless of the type of technology in use. Within the cloud operating environment, there is a high degree of complexity, because varying laws and regulations may apply within different countries, states and localities.

External requirements help to support the reasons why policies, standards and controls must be established within the organization, and they dictate what assurances may be required from the CSP. For example, US federal government entities that leverage a cloud provider not only have specific control objectives to achieve for their responsibilities, but also require the CSP to provide evidence of adherence to the Federal Risk and Authorization Management Program (FedRAMP) requirements. An organization that outsources operations to a CSP remains accountable for adherence to regulations.

Similarly, there are several standards available that can guide an organization cloud controls program. These standards are typically established to provide uniformity and guidance on techniques for addressing risk through controls. Although not always as enforceable as regulatory requirements, these standards can help an organization demonstrate that it is using sound practices to support its cloud operations.

Laws, regulations and standards, and many organization operations are constantly changing. Organizations should take care to maintain an understanding of the relevant laws, regulations and standards. What was once deemed an irrelevant standard, law or regulation may become relevant if the organization acquires an entity or otherwise expands operations. Additionally, as laws, regulations and standards are updated to align with the evolving environment, there may be new requirements or implementation considerations for an organization.

2.6.1 Legal and Regulatory Requirements

Some requirements have their origins in regulations or laws required by specific industries. These requirements commonly exist to stop personal information from falling victim to theft or fraud. For example:

- If an organization deals with health sector providers that share confidential patient data in the United States, it needs to comply with the Health Insurance Portability and Accountability Act (HIPAA) and Health Information Technology for Economic and Clinical Health (HITECH) Act. HIPAA sets a national standard and requirements to protect health and medical records and other forms of protected health information. The HITECH Act defines electronic health records and establishes breach notification requirements, among other things.

- In the US financial services industry, companies must comply with the Gramm-Leach-Bliley Act (GLBA) in developing and maintaining information security programs to protect customers' private information. GLBA requires financial services companies to maintain a written information program that includes administrative, operational and technical safeguards.

Location-based laws occur at a variety of levels, from national laws to state and local laws. For example, within the United States, several states have stand-alone security and privacy laws that apply to operations within their borders or to consumer interactions within those states. Similar granularity occurs elsewhere in the world, such as in Europe and Asia.

The European Union industry-agnostic General Data Protection Regulation (GDPR), a primary regulation with broad reach, is centered on the concept of protecting the individual's right to privacy. It applies to any organization that processes information about persons resident in the EU, even if the organization is not physically located there. GDPR is not prescriptive about technology but requires organizations to take adequate steps to protect personal data. Organizations and their service providers must be aware of its potential applicability and implications.

Requirements associated with laws and regulations commonly pass through an organization to relevant service providers; relevant service providers (such as CSPs) must also demonstrate adherence. These requirements are typically codified in contractual arrangements, specifying what the cloud service provider must adhere to on behalf of its customer. For example, many US government contracts require that cloud services comply with the Federal Information Security Management Act (FISMA).[35] The FedRAMP program was designed to support FISMA adherence by providing a standardized approach for security in the cloud, including a baseline of controls with various implementation levels based on risk, standard assessment activities and authorized third-party assessment firms. Examples of other regulations that might apply follow:

- The Sarbanes-Oxley Act, which sets requirements for publicly traded companies in the United States;[36] the Financial Instruments and Exchange Act (i.e., J-SOX), which sets similar requirements in Japan[37]
- Privacy rules, including the Canada Personal Information Protection and Electronic Documents Act (PIPEDA),[38] the Australian Privacy Principles (APP),[39] the New York Department of Financial Services Cybersecurity Regulation[40] and the California Consumer Privacy Act[41]
- The Family Educational Rights and Privacy Act (FERPA), which is related to educational records[42]
- The EU Directive on security of network and information systems (NIS Directive)[43]
- The US Clarifying Lawful Overseas Use of Data (CLOUD) Act,[44] which was designed to speed access to data stored by US-based global companies to support criminal investigations

Organizations should have a clear understanding of applicable laws and regulations impacting their business operations, including:

[35] govinfo, "Public Law 107-347," 17 December 2002, www.govinfo.gov/content/pkg/PLAW-107publ347/pdf/PLAW-107publ347.pdf

[36] govinfo, "Public Law 107-204," 30 July 2002, www.govinfo.gov/content/pkg/PLAW-107publ204/pdf/PLAW-107publ204.pdf

[37] Financial Services Agency, "Financial Instruments and Exchange Act," 7 June 2006, www.fsa.go.jp/en/policy/fiel/

[38] Office of the Privacy Commissioner of Canada, "The Personal Information Protection and Electronic Documents Act (PIPEDA), www.priv.gc.ca/en/privacy-topics/privacy-laws-in-canada/the-personal-information-protection-and-electronic-documents-act-pipeda/

[39] Office of the Australian Information Commissioner, "Australian Privacy Principles," www.oaic.gov.au/privacy/australian-privacy-principles/

[40] De Groot, J.; "What is the NYDFS Cybersecurity Regulation? A Cybersecurity Compliance Requirement for Financial Institutions," Data Insider, 24 October 2019, https://digitalguardian.com/blog/what-nydfs-cybersecurity-regulation-new-cybersecurity-compliance-requirement-financial

[41] State of California Department of Justice, "California Consumer Privacy Act (CCPA)," https://oag.ca.gov/privacy/ccpa

[42] US Department of Education, "Family Educational Rights and Privacy Act (FERPA)," www2.ed.gov/policy/gen/guid/fpco/ferpa/index.html

[43] European Commission, "The Directive on security of network and information systems (NIS Directive)," https://ec.europa.eu/digital-single-market/en/network-and-information-security-nis-directive

[44] US Department of Justice, "Promoting Public Safety, Privacy, and the Rule of Law Around the World: The Purpose and Impact of the CLOUD Act," April 2019, www.justice.gov/opa/press-release/file/1153446/download

- Impacts to control requirements
- Extent of applicability to service providers
- Processes to maintain the inventory

2.6.2 Cloud Standards and Best Practices

Cloud Standards Background

Cloud computing is a new technical and organizational IT paradigm that requires a new approach to certain aspects of information technology, networking and security, including new tools, best practices, frameworks, and, sometimes, new standards.

During the early stages of development and adoption of cloud technologies, there was a long debate on the lack of standards dedicated to the cloud. Some key organizations, such as the European Telecommunications Standards Institute (ETSI), the European Commission and NIST, supported research to take stock of existing standards and analyze gaps that would prevent the safe adoption of cloud computing.

Their studies consistently cautioned that cloud computing is not completely new and that many standards used for cloud computing are not cloud specific. In the future, these standards may continue to evolve to better reflect cloud computing scenarios.[45] This cycle of iteration and new standard release constantly repeats itself, with updates released frequently from a number of sources.

The biggest gaps in the cloud standardization landscape were evident in the areas of interoperability and portability, service level agreements, and, in part, security and privacy. During the period from 2009 to 2019, several standards, best practices and guidelines were published by international and national bodies, and some other sectoral organizations, communities and forums.

Notable examples of standard development organizations (SDOs), agencies and research organizations developing cloud-relevant standards follow:

- ISO/IEC, ITU-T, IEEE, among international SDOs
- NIST (US), European Union Agency for Cybersecurity (ENISA) and ETSI (Europe), Federal Office for Information Security (BSI) (Germany), IMDA (Singapore), National Cybersecurity Agency of France (ANSSI) (France), among the national or regional institutions
- Cloud Security Alliance (CSA), DMTF, OASIS, OWASP, FIDO Alliance, OGF, IETF, among international organizations, community, forums and alliances
- PCI Security Standards Council (PCI SSC), among sector-specific organizations

ISO/IEC, NIST and CSA, over the last decade, produced the most comprehensive sets of work on cloud definition, architecture, security, privacy and interoperability.

ISO/IEC developed a series of standards: cloud definitions (e.g., ISO/IEC 17788[46]), cloud architecture (ISO/IEC 17789[47]), cloud security (ISO/IEC 27017[48]), cloud privacy (ISO/IEC 27018[49]), Cloud SLAs (ISO/IEC 19086 Parts

[45] ETSI, "Cloud Standards Coordination Final Report," November 2013, https://ec.europa.eu/digital-single-market/en/news/cloud-standards-coordination-final-report

[46] ISO, ISO/IEC 17788:2014 *Information technology – Cloud computing – Overview and vocabulary*, https://www.iso.org/standard/60544.html

[47] ISO, ISO/IEC 17789:2014 *Information technology – Cloud computing – Reference architecture*, https://www.iso.org/standard/60545.html

[48] ISO, ISO/IEC 27017:2015 *Information technology – Security techniques – Code of practice for information security controls based on ISO/IEC 27002 for cloud services*, https://www.iso.org/standard/43757.html

[49] ISO, ISO/IEC 27018:2014 *Information technology – Security techniques – Code of practice for protection of Personally Identifiable Information (PII) in public clouds acting as PII processors*, https://www.iso.org/standard/61498.html

1[50], 2[51], 3[52] and 4[53]), virtualization (ISO/IEC 21878[54] and ISO/IEC 19099[55]).

NIST released the first and arguably the most widely used definition of cloud computing (Special Publication 145[56]), a cloud computing reference architecture (SP 500-292[57]) and security architecture (SP 500-299[58]), an evaluation approach for cloud computing (SP 500-322[59]), and security and privacy controls (SP 800-53[60]). SP 800-53 started as a general-purpose framework but evolved to become cloud relevant.

Since its formation in 2009, the Cloud Security Alliance generated several best practices, guidelines, frameworks and recommendations for cloud computing. Among the most relevant and widely used are:

- Cloud Controls Matrix (CCM[61])
- Guidance for Critical Areas of focus in cloud computing[62]
- Consensus Assessment Initiative Questionnaire (CAIQ[63])
- Open Certification Framework (OCF[64])
- Privacy Level Agreement (PLA[65])
- GDPR Code of Conduct[66]
- Cloud Enterprise Architecture[67] (on which the NIST cloud security architecture is based)
- Software Defined Perimeter (SDP[68])

Standards Relevance for Compliance

Standards can serve as an additional input to an organization cloud controls program. Compared to laws and regulations, standards tend to apply more granularity in how to achieve control objectives and address risk. Organizations and cloud service providers leverage the standards for multiple purposes, including the following:

[50] ISO, ISO/IEC 19086-1:2016 *Information technology – Cloud computing – Service level agreement (SLA) framework – Part 1: Overview and concepts*, https://www.iso.org/standard/67545.html

[51] ISO, ISO/IEC 19086-2:2018 *Cloud computing – Service level agreement (SLA) framework – Part 2: Metric model*, https://www.iso.org/standard/67546.html

[52] ISO, ISO/IEC 19086-3:2017 *Information technology – Cloud computing – Service level agreement (SLA) framework – Part 3: Core conformance requirements*, https://www.iso.org/obp/ui/#iso:std:iso-iec:19086:-3:ed-1:v1:en

[53] ISO, ISO/IEC 19086-4:2019 *Cloud computing – Service level agreement (SLA) framework – Part 4: Components of security and of protection of PII*, https://www.iso.org/standard/68242.html

[54] ISO, ISO/IEC 21878:2018(en) *Information technology – Security techniques – Security guidelines for design and implementation of virtualized servers*, https://www.iso.org/standard/72029.html

[55] ISO, ISO/IEC 19099:2014 *Information technology – Virtualization Management Specification*, https://www.iso.org/standard/63962.html

[56] Computer Security Resource Center, "The NIST Definition of Cloud Computing," September 2011, https://csrc.nist.gov/publications/detail/sp/800-145/final

[57] Liu, F.; J. Tong; J. Mao; R. Bohn; J. Messina; L. Badger; D. Leaf; "NIST Cloud Computing Reference Architecture," National Institute of Standards and Technology, September 2011, https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication500-292.pdf

[58] Computer Security Resource Center, "NIST Cloud Computing Security Reference Architecture," May 2013, https://csrc.nist.gov/publications/detail/sp/500-299/draft

[59] Simmon, E.; "Evaluation of Cloud Computing Services Based on NIST SP 800-145," February 2018, https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.500-322.pdf

[60] Computer Security Resource Center, "Security and Privacy Controls for Information Systems and Organizations," September 2020, https://csrc.nist.gov/publications/detail/sp/800-53/rev-5/final

[61] Cloud Security Alliance, "Cloud Controls Matrix (CCM)," https://cloudsecurityalliance.org/research/cloud-controls-matrix/

[62] Cloud Security Alliance, "CSA Security Guidance for Critical Areas of Focus in Cloud Computing v4.0," https://cloudsecurityalliance.org/research/guidance/

[63] Cloud Security Alliance, "Consensus Assessment Initiative Questionnaire (CAIQ) v3.1," 1 April 2020, https://cloudsecurityalliance.org/artifacts/consensus-assessments-initiative-questionnaire-v3-1/

[64] Cloud Security Alliance, "Working Group – Open Certification Framework," https://cloudsecurityalliance.org/research/working-groups/open-certification/

[65] Cloud Security Alliance, "Working Group – Privacy Level Agreement," https://cloudsecurityalliance.org/research/working-groups/privacy-level-agreement/

[66] Cloud Security Alliance, "CSA CoC for GDPR Compliance," https://cloudsecurityalliance.org/privacy/gdpr/code-of-conduct/

[67] Cloud Security Alliance, "Working Group – Enterprise Architecture," https://cloudsecurityalliance.org/research/working-groups/enterprise-architecture/

[68] Cloud Security Alliance, "Working Group – Software Defined Perimeter," https://cloudsecurityalliance.org/research/working-groups/software-defined-perimeter/

- Defining their high-level compliance framework
- Speaking a common language
- Identifying potential controls to address known risk
- Driving uniformity in control techniques
- Demonstrating sufficiency of practices to address known risk

For cloud customers, a key component of establishing a cloud controls program is clearly identifying the authoritative sources to use. Organizations usually identify several authoritative sources, although in some cases an organization opts to identify one standard or framework as the primary guiding principle.

For a CSP, it is important to define its strategy and identify its target customer base. This approach will provide further clarity as to what standards to leverage in building a security program and which third-party assessment vehicles to obtain.

Tip: Like with laws and regulations, standards are routinely updated and their applicability to an organization evolves in alignment with organization changes. The cloud controls program should be structured and governed so that the framework is able to dynamically adjust in response to any relevant changes. Organizations usually maintain an inventory of laws, regulations and standards mapped to requirements to facilitate this dynamic update process.

Cloud Security Frameworks

General-purpose information security frameworks, such as ISO/IEC 27001 and 27002 (see section 2.10 for additional details on ISO/IEC 27001), offer a good foundation to address key security requirements. They typically contain controls about the definition, implementation, enforcement and review of policies for incident management, or about the testing of certain procedures, or about the enforcement of key principles, such as segregation of duties or least privileges. Those controls are relevant to cloud computing too.

However, there are situations in which generic information security frameworks fall short in addressing the specificity of cloud computing. Certain areas—such as the shared responsibility model, supply chain accountability, workload segregation, network segregation, virtualization, auditing, business continuity testing and compliance—deserve a new approach in the cloud.

To fill the gaps existing in information security frameworks, several organizations over time decided to develop cloud-specific or cloud-relevant frameworks. ENISA made the first attempt at the end of 2009 when it published the Information Assurance Framework (IAF). In 2010, the Cloud Security Alliance released the Cloud Controls Matrix.

Commonly Adopted Frameworks

There are many security frameworks that provide foundations to help efficiently and effectively evaluate, assess, select, deploy, configure, manage, optimize and secure cloud services. The frameworks described here are just some of those frequently referenced. These frameworks are often reused and extended, and the same experts often contribute to multiple frameworks, providing a level of consistency and familiarity.

CCM

The Cloud Security Alliance Cloud Controls Matrix (CCM) is a set of security controls specifically designed to provide fundamental security principles to guide CSPs and to assist prospective cloud customers in assessing the overall security risk of a cloud provider. The CSA CCM provides a controls framework that provides a detailed understanding of security concepts and principles aligned to the Cloud Security Alliance Security Guidance.

The foundations of the Cloud Security Alliance Controls Matrix rest on its customized relationship to other industry-accepted security standards, regulations and controls frameworks, such as ISO/IEC 27001/27002/27017/27018, NIST, PCI, BSI, ENISA, ISACA COBIT 2019 and NERC CIP. The CCM augments or provides internal control direction for service organization control reports, attestations that CSPs provide. Chapter 3 is dedicated to the CCM.

ISO/IEC 27017

The International Organization for Standardization (ISO) publishes a series of information security standards that can support a cloud controls compliance program. The ISO/IEC 27001 standard is broadly applicable to any organization, because it provides a specification for an Information Security Management System (ISMS). ISO/IEC 27002 describes controls that can be put in place to adhere to the ISO/IEC 27001 standard. Further building on these foundational pieces, ISO published ISO/IEC 27017, which provides guidance on the information security aspects of cloud computing, recommending and assisting with the implementation of cloud-specific information security controls supplementing the guidance in ISO/IEC 27002.

ISO/IEC 27017 is a code of practice meant to provide cloud-related information on how to implement the security controls included in ISO/IEC 27002 and additional controls specific to the cloud computing context. The seven new cloud-specific controls focus on areas, such as shared roles and responsibilities, management of cloud assets, and segregation of cloud environment networks.

The standard advises cloud customers and CSPs, with the primary guidance laid out side-by-side in each section of the standard documentation.

ISO/IEC 27018 and ISO/IEC 27701

ISO/IEC 27018 is a code of practice for protecting personally identifiable information (PII) in public clouds acting as PII processors. ISO/IEC 27018 also can be relevant to organizations acting as PII controllers (which might be subject to additional requirements). It includes control objectives, controls and guidelines for implementing measures to protect PII in line with the privacy principles in ISO/IEC 29100 for the public cloud computing environment.

Like ISO/IEC 27017, ISO/IEC 27018 specifies "guidelines based on ISO/IEC 27002, taking into consideration the regulatory requirements for the protection of PII which can be applicable within the context of the information security risk environment(s) of a provider of public cloud services."[69]

The target audience for this standard includes public and private companies, government entities and not-for-profit organizations of any type and size.

In 2019, the ISO Standardization Subcommittee (SC) 27 released ISO/IEC 27701:2019 *Security techniques – Extension to ISO/IEC 27001 and ISO/IEC 27002 for privacy information management – Requirements and guidelines*.[70] This new standard is a privacy extension to ISO/IEC 27001. It is meant to provide organizations that have already implemented an information security management system (ISMS) with additional guidance for the protection of privacy. The goal is to support organizations interested in establishing a privacy information management system (PIMS).

[69] ISO, ISO/IEC 27018:2019 *Information technology – Security techniques – Code of practice for protection of personally identifiable information (PII) in public clouds acting as PII processors*, www.iso.org/standard/76559.html

[70] ISO, ISO/IEC 27701:2019 *Security techniques – Extension to ISO/IEC 27001 and ISO/IEC 27002 for privacy information management – Requirements and guidelines*, https://www.iso.org/standard/71670.html

NIST SP 800-53

NIST SP 800-53 Security and Privacy Controls for Information Systems and Organizations "provides a catalog of security and privacy controls for federal information systems and organizations to protect organizational operations and assets, individuals, other organizations, and the Nation from a diverse set of threats and risks, including hostile attacks, natural disasters, structural failures, human errors, and privacy risks."[71] NIST SP 800-53 was not created as a cloud-specific standard, but it has evolved to become cloud relevant. The most recent publication (Rev 5) was published in September 2020. According to NIST, Rev 5 includes the following changes compared to Rev 4:

- **Consolidates the control catalog**—Information security and privacy controls are now integrated into a seamless, consolidated control catalog for information systems and organizations.
- **Integrates supply chain risk management**—Rev. 5 establishes a new supply chain risk management (SCRM) control family and integrates SCRM aspects throughout the catalog.
- **Adds new state-of-the-practice controls**—These are based on the latest threat intelligence and cyberattack data (e.g., controls to support cyberresilience, secure systems design, security and privacy governance, and accountability).
- **Makes controls outcome-based**—Rev. 5 accomplishes this by removing the entity responsible for satisfying the control (i.e., information system or organization) from the control statement.
- **Improves descriptions of content relationships**—Rev. 5 clarifies the relationship between requirements and controls, and the relationship between security and privacy controls.
- **Separates the control selection processes from the controls**—This allows the controls to be used by different communities of interest, including systems engineers, security architects, software developers, enterprise architects, systems security and privacy engineers, and mission or business owners.
- **Transfers control baselines and tailoring guidance to NIST SP 800-53B**—This content has moved to the new (draft) Control Baselines for Information Systems and Organizations.

BSI C5

The BSI C5[72] standard is the Cloud Computing Compliance Controls Catalog (C5) for the German government-backed attestation scheme introduced by the Federal Office for Information Security (BSI) to help organizations demonstrate operational security against common cyberattacks within the context of the German government Security Recommendations for Cloud Providers.

The basic criteria of the BSI C5 standard reflect the minimum level of information security that a cloud service must offer when cloud customers use it to process information that has a normal need for protection. The minimum scope of an audit is defined by this criteria catalog. It is up to the cloud customer to assess the extent to which the basic criteria adequately reflect the protection needs of its information in its individual use case.

For cloud customers whose information needs greater protection, additional criteria provide a starting point for conducting this assessment. CSPs may expand an audit based on the basic criteria by including the additional criteria for customers with higher protection needs.

ENISA Cloud Computing IAF

The European Union Agency for Cybersecurity (ENISA) is an EU body whose mission is to achieve a high, common level of cybersecurity across the European Union. In 2009, ENISA published the Cloud Computing Information

[71] National Institute of Standards and Technology, "Security and Privacy Controls for Information Systems and Organizations," March 2020, https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r5-draft.pdf

[72] BSI, "C5 Introduction," www.bsi.bund.de/EN/Topics/CloudComputing/Compliance_Criteria_Catalogue/C5_Introduction/C5_Introduction_node.html

Assurance Framework[73] (IAF), which provides a set of assurance criteria designed to assess cloud-related risks, compare cloud offerings, obtain assurance from providers and reduce the assurance burden on cloud providers. The IAF was included as a component of ENISA Cloud Computing Risk Assessment,[74] a study published to analyze the potential risk associated with adopting cloud computing. The IAF highlights the key questions an organization should ask when building a cloud security program to manage the fundamental risk identified in the report. The IAF has not been updated since 2009; although the principles it embodies are still relevant, some control questions might be outdated and a few gaps can be present.

Others

Many security frameworks and compliance control sets exist, each focusing on a different targeted audience. The following sections present a subset of groups and communities.

PCI DSS

The Payment Card Industry Data Security Standard (PCI DSS) was developed by the PCI Security Standards Council (PCI SSC) to establish consistent security practices for securing cardholder data. The standard is applicable to a variety of organizations that process or store credit card payment information. The standard includes a series of technical and operational requirements, related implementation guidance, and testing procedures for assessment firms. PCI-DSS can inform the control requirement of an organization and the requirements that it may have of its service provider. Like with other standards, a cloud service provider can obtain an attestation of compliance from a qualified assessment firm to demonstrate to its customers its adherence to the requirements.

PCI SSC Cloud Computing Guidelines

Through the cloud computing guidelines,[75] the PCI SSC provides guidance on how to implement PCI-DSS requirements in a cloud-based environment.

Center for Internet Security (CIS)—Controls and Benchmark

The CIS published useful tools to support cybersecurity, particularly the CIS Controls and CIS Benchmarks. The CIS Controls[76] include a set of technical controls generically applicable to information systems and networks.

The CIS Benchmarks include security benchmarks that are published for a variety of technologies and platforms. CIS provides cloud benchmarks[77] for some major cloud service providers, such as AWS®, Microsoft® Azure and Google Cloud™ enterprise services.

ANSSI—SecNumCloud

The National Cybersecurity Agency of France (ANSSI) published the SecNumCloud—Requirements Repository—Essential level – v.3, and the Advanced Requirements.

The two documents contain requirements for secure cloud computing and cloud service providers applicable to organizations that provide services to administrative authorities.

[73] European Union Agency for Cybersecurity, "Cloud Computing Information Assurance Framework," 20 November 2009, www.enisa.europa.eu/publications/cloud-computing-information-assurance-framework

[74] European Union Agency for Cybersecurity, "Cloud Computing Risk Assessment," 20 November 2009, www.enisa.europa.eu/publications/cloud-computing-risk-assessment

[75] PCI Security Standards Council, "Information Supplement: PCI SSC Cloud Computing Guidelines," April 2018, www.pcisecuritystandards.org/pdfs/PCI_SSC_Cloud_Guidelines_v3.pdf

[76] Center for Internet Security, "CIS Controls," www.cisecurity.org/controls/

[77] Center for Internet Security, "CIS Benchmarks," www.cisecurity.org/cis-benchmarks/

EU-SEC

The EU-SEC consortium developed a framework for defining guidelines, procedures and tools to foster mutual recognition for existing cloud security certifications. For a part of this work, it created a repository of security and privacy controls that includes a catalog of controls derived from CCM and the gap analysis between CCM and other cloud-relevant standards. It can be considered an extension of the CCM; the CSA used the results of the EU-SEC work to develop Addenda to CCM.

The project was funded by the European Commission under the H2020 Framework Program, and the CSA was a member of the consortium and the scientific coordinator of the project.

Resources are freely available at the EU-SEC project website.[78]

Relationship Between Existing Frameworks

Cloud-relevant or cloud-specific frameworks function as an extension of general-purpose information security frameworks, and there is consequently a strong relationship between them. The way that organizations select and use those frameworks largely depends on factors such as the following:

- The maturity of the organization security compliance program
- Whether the organization is born in the cloud or is moving to the cloud
- Business requirements
- Certification and attestation requirements

An organization that already has a security compliance program in place is likely to consider ways to extend the existing set of controls by integrating cloud-specific controls to fill potential gaps.

An organization that is mainly cloud-based might adopt a cloud-native control framework, since that would cover all the necessary security requirements of best practices.

An organization with a complex and mature approach to security and compliance might develop a tailored control framework involving several control frameworks.

An organization operating in several business sectors and countries may adopt and adhere to the requirements of several frameworks to present its compliance posture in the easiest way possible to a customer or authority of reference.

Regardless of the approach and framework selected, it is important for the cloud customer to map the organizational security requirements with the chosen framework—or frameworks, in the case of the multiframework approach—to clearly understand the relationships between them.

There is considerable overlap between the requirements included in the various control frameworks. Often, different frameworks include controls that include semantic differences, but that are substantially the same from a practical point of view. For example, according to the results of a study[79] conducted by the EU-SEC consortium, the controls included in the CCM cover approximately 85 percent of the requirements included in ISO/IEC 27001/2/17/18, C5, AICPA TSC and other standards.

[78] EU Security Certification, "The European Security Certification Framework (EU-SEC), www.sec-cert.eu/
[79] EU Security Certification, "European Security Certification Framework D 1.2 Security and Privacy Requirements and Controls," 28 May 2019, https://cdn0.scrvt.com/fokus/5c1e85c636d71d7c/c559f8e809fc/EU-SEC-D1.2-Security-and-privacy-requirements-and-controls-V1.4.pdf

2.7 Defining Controls

2.7.1 Controls

ISO defines a control as a measure that modifies a risk.[80] According to ISACA, internal controls include the policies, standards, procedures and other organizational structures that are designed to provide reasonable assurance that business objectives will be achieved, and undesired events will be prevented, detected and corrected.[81] For example, installing antivirus software is a technical control that aims to reduce a broad set of risk related to malware infections.

A control is any measure that can help safeguard the security of physical property, information, computer systems or other assets; or any countermeasure that can help avoid, detect, counteract or minimize security risk to physical property, information, computer systems or other assets. Controls can be classified by several criteria, but there are three main categories:

- **Preventive**—Measures reducing or eliminating the likelihood or impact of a risk. Note that ISO redefined preventive to mean actions to address risk and opportunities.[82]
- **Detective**—Measures that detect and characterize events[83] that lead to the realization of a risk
- **Corrective**—Measures designed for risk impact reduction after detection

An example of a preventive control is an antivirus software installation. A detective control can include an intrusion detection system that monitors network traffic for suspicious activities. A corrective control can be the system spraying fire suppressant into a data center room in case of a fire, or an AWS Lambda function closing an S3 bucket that is open to the Internet.

Controls are also often categorized[84, 85] according to their nature, as follows:

- **Technical controls (logical)**—Software, hardware and related technologies that are used to protect assets, including antivirus measures, encryption and network filtering
- **Administrative/organizational controls**—Policies, regulations, standards and procedures that constrain organization management and protection of its assets, including measures such as training, background checks on personnel, and supply chain management[86]
- **Physical controls**—Devices that protect access and integrity of the environment where computing resources and personnel are located, including locks on doors, guards, and security cameras

High-level or descriptive control languages may express requirements covering more than one of the categories described. See chapter 4 for additional details about controls categories and examples.

[80] ISO/IEC 27000:2014

[81] ISACA, “Glossary,” www.isaca.org/Pages/Glossary.aspx?tid=1506&char=I

[82] A security system by its very nature is a preventive tool. Therefore, Preventive Action has been removed from ISO directives and redefined as Risk-based ... Plan Actions to address Risks and Opportunities.

[83] An event is defined as “occurrence or change of a particular set of circumstances” (ISO Guide 73:2009).

[84] Red Hat, “Red Hat Enterprise Linux 4 Security Guide,” https://access.redhat.com/documentation/en-us/red_hat_enterprise_linux/4/pdf/security_guide/Red_Hat_Enterprise_Linux-4-Security_Guide-en-US.pdf Section 1.2

[85] European Union Agency for Cybersecurity, “Guidelines for SMEs on the security of personal data processing,” 27 January 2017, www.enisa.europa.eu/publications/guidelines-for-smes-on-the-security-of-personal-data-processing

[86] Some sources add a fourth category, such as “regulatory controls,” distinct from administrative controls.

Figure 2.9 shows three examples of controls from the CSA Cloud Controls Matrix version 3.0.1, each covering a different control category.

Figure 2.9—CSA Cloud Control Categories

Category	Example
Technical control	Data input and output integrity routines (i.e., reconciliation and edit checks) shall be implemented for application interfaces and databases to prevent manual or systematic processing errors, corruption of data, or misuse. (CCM AIS-03)
Administrative/ Organizational control	Managers are responsible for maintaining awareness of, and complying with, security policies, procedures, and standards that are relevant to their area of responsibility. (CCM GRM-03)
Physical control	Ingress and egress to secure areas shall be constrained and monitored by physical access control mechanisms to ensure that only authorized personnel are allowed access. (CCM DCS-7)

Controls such as the one found in control frameworks like the CSA CCMv3.0.1 are high-level controls that combine several control categories. Consider, for example, CCM EKM-04.

> Example: CCM V3.0.1 EKM-04
>
> Platform- and data-appropriate encryption (e.g., AES-256) in open/validated formats and standard algorithms shall be required. Keys shall not be stored in the cloud (i.e., at the cloud provider in question), but maintained by the cloud consumer or trusted key management provider. Key management and key usage shall be separated duties.

Although this is globally a preventive control, its content can be divided into two parts. The first sentence describes a technical control: Platform and data appropriate encryption (e.g., AES-256) in open/validated formats and standard algorithms shall be required. The second and third sentences are mainly organizational matters.

2.7.2 Control Objectives

Some control frameworks group controls together under a common control objective. A control objective is a statement describing what is to be achieved as a result of implementing controls.[87]

Control objectives are more abstract, or high-level, statements than controls. In practice, not all audit frameworks make the distinction between controls and control objectives. For example, ISO/IEC 27002 provides a catalog of control objectives, with each associated with one or more controls; NIST SP 800-53 refers only to controls (and control enhancements).

Other sources use a different terminology altogether to refer to control objectives. For example, the AICPA[88] prefers the term trust services criteria in the context of attestations or consulting engagements,[89] to refer to control objectives.

The following examples from the annex of ISO 27001:2005 illustrate the difference between a control and a control objective:

- Control objective (A.9.1 Secure areas): To prevent unauthorized physical access, damage and interference to the organization's premises and information.
- Control (A.9.1.4): Physical protection against damage from fire, flood, earthquake, explosion, civil unrest, and other forms of natural or man-made disaster shall be designed and applied.

[87] From ISO/IEC 27000:20014

[88] American Institute of Certified Public Accountants, www.aicpa.org/

[89] AICPA, "TSP Section 100 – 2017 Trust Services Criteria for Security, Availability, Processing Integrity, Confidentiality, and Privacy," revised March 2020, www.aicpa.org/content/dam/aicpa/interestareas/frc/assuranceadvisoryservices/downloadabledocuments/trust-services-criteria.pdf

2.7.3 Designing and Selecting Controls

The design and selection of controls in an information system is not an activity that stands on its own, but is part of the broader process of risk management. By definition, controls are countermeasures that address a risk. Selected and implemented controls should be regularly reviewed for effectiveness. Organizations may alter or replace ineffective controls to meet their objectives.

After risk is identified, priority goes to the mitigation of the highest risk, through the selection of appropriate controls. Although risk should be treated according to its level of severity (high, medium, low), it may be possible to treat more than one risk severity at the same time. In some cases, it is acceptable not to treat a risk condition, simply because it is sufficiently low, or because the cost of treating the risk far outweighs the cost of the realization of the risk itself. Each organization has a different risk tolerance. Best practices encourage documenting and reporting the management decision to accept a risk to the executive team, board of directors, or audit committee of the board of directors.

Controls can be designed and selected to reduce the impact or likelihood of a risk, or both. For example, using specific cryptographic hashing techniques[90] to protect user passwords is a way to reduce the impact of a breach of confidentiality of the user database containing the passwords, because those passwords are not exploitable by the adversary in most cases. Likewise, removing all administrative remote access mechanisms to a virtual machine (SSH, remote web console, etc.) reduces the likelihood of a remote compromise. A just-in-time access approach reduces the time access is available to the machine, thus reducing the window of opportunity for unauthorized access.

In the context of cloud computing and the use of third-party services, it is mandatory to take into account the controls that the third party already offers through security tools, APIs and configuration options. However, those controls may not be sufficient to address the identified risk, and it may be necessary to complement or supplement them with additional controls.

Risk management is an ongoing process. After controls are implemented, it is important to continuously monitor and assess their effectiveness. Gaps and shortcomings may be identified, requiring controls to be modified, removed or added. Moreover, as the organization grows and threats evolve or change, new risk emerges, requiring the selection and implementation of new controls.

In a public cloud with a DevOps IT environment, the traditional risk management process is challenged, and it might be unable to keep up with the pace of daily changes and deployments to production. In such cases, the DevOps team is often responsible for mitigating or managing risk daily. The risk management team checks afterward if DevOps followed processes and procedures and generated the required evidence. The internal audit function typically determines if the approach used is mature and effective, by assessing whether the action taken addresses the risk and supports the organization objectives, and by evaluating the balance between preventive, detective and corrective controls.

2.7.4 Control Frameworks

Risk management involves the identification of risk followed by the selection of controls that mitigate them. While every organization faces different risk and uses different information systems, there are many commonalities in the way risk must be addressed, leading ultimately to many commonalities in control objectives. All organizations, including the smallest, need clearly defined information security management practices and governance. Based on the experience gained by security practitioners over the years, this led standardization organizations to develop comprehensive and logically structured catalogs of controls that can be reused and applied across a wide range of organizations. These catalogs are often referred to as control frameworks.

[90] For example, the PBKDF2 (Password-Based Key Derivation Function 2)

The use of a standardized control framework does not mean that an organization implements all controls in the framework. Some proposed controls might not be relevant and are ignored. In some cases, organizations might add specific controls that may be missing from the framework. In other cases, organizations adopt alternative controls that achieve the same results as the ones that the framework suggests. The selection of a control must be grounded in the results of a risk analysis, and a control framework only provides a well-defined set of measures for risk treatment.

Using a standardized control framework has many benefits, including the following:

- It anchors the implementation of security controls into a well-known reference system and language recognized by many practitioners across an industry.
- It supports auditors' efforts to assess an information system based on a well-defined set of controls, enabling the following actions:
    - Comparisons of cloud security features across different time periods or different cloud services
    - Gap analyses of what a CSP offers versus what the customer requires
    - Benchmarking of competing cloud services (e.g., in the context of procurement)
- It enables the building of trusted certifications schemes by ensuring that information systems are assessed with a comparable set of criteria.

Control frameworks are usually organized in a hierarchical manner. First, controls are grouped into domains designed to help understand their broad scope. Domains may also be called families or chapters. Domains may help divide audit and compliance efforts across organizations. For instance, controls grouped together under the physical and environmental security domain and those grouped under the configuration management domain might be overseen by different managers and auditors.

Some control frameworks break down each domain into a set of control objectives, and then break down each control objective into one or more controls. Other control frameworks adopt a simpler hierarchy that simply groups controls into domains. Most control frameworks provide implementation guidelines for each control.

Most control frameworks assign a unique control identifier to each control—i.e., a short name or number that facilitates cross-referencing.

Examples

ISO/IEC 27002:2013 is a generic control framework for information security management systems. The standard defines 14 domains, each divided into control objectives and then further subdivided into controls. A total of 35 control objectives cover 114 individual controls.

NIST SP 800-53 is a generic control framework designed for US federal information systems. The standard defines 18 control families covering a total of 231 controls (**figure 2.10**). Controls are classified into three baselines—low, moderate and high—each designed to address a particular risk level selected through a preliminary risk analysis.

Figure 2.10—NIST SP 800-53r4 Control Families

ID	FAMILY	ID	FAMILY
AC	Access Control	**MP**	Media Protection
AT	Awareness and Training	**PE**	Physical and Environmental Protection
AU	Audit and Accountability	**PL**	Planning
CA	Security Assessment and Authorization	**PS**	Personal Security
CM	Configuration Management	**RA**	Risk Assessment
CP	Contingency Planning	**SA**	System and Services Acquisition
IA	Identification and Authentication	**SC**	System and Communications Protection
IR	Incident Response	**SI**	System and Information Integrity
MA	Maintenance	**PM**	Program Management

The CSA Cloud Controls Matrix v3.0.1 is a control framework designed to specifically address cloud services and cloud threats. The CCM is divided into 16 domains covering 133 control objectives (**figure 2.11**).

Figure 2.11—CSA CCM v3.0.1 Control Domains

AIS	Application & Interface Security	HRS	Human Resources Security
AAC	Audit Assurance & Compliance	IAM	Identity and Access Management
BCR	Business Continuity Mgmt & Op Resilience	IVS	Infrastructure & Virtualization
CCC	Change Control & Configuration Management	IPY	Interoperability & Portability
DSI	Data Security & Information Lifecycle Mgmt	MOS	Mobile Security
DCS	Datacenter Security	SEF	Sec. Incident Mgmt, E-Disc & Cloud Forensics
EKM	Encryption & Key Management	STA	Supply Chain Mgmt, Transparency & Accountability
GRM	Governance & Risk Management	TVM	Threat & Vulnerability Management

Source: Cloud Security Alliance, Cloud Controls Matrix (CCM) Working Group

See chapter 3 for more details on the CSA Cloud Controls Matrix.

2.7.5 Mapping Controls across Control Frameworks

Many organizations need to show compliance with more than one control framework to accommodate customers and comply with regulations across industries and geographical regions. It is important to understand which compliance frameworks are relevant to the cloud services that an organization uses or offers.

Organizations assessing compliance with more than one framework have two choices: assess compliance with each framework separately, starting each time from scratch; or look for commonalities between frameworks and avoid duplication of work where possible.

There are many common controls across different frameworks. The wording may differ, but the objectives are often the same. Therefore, the second approach makes a lot of sense. Organizations are encouraged to maintain mappings between the different control frameworks that they review.

See chapter 3 for more information on cross-mapping of control frameworks.

2.7.6 Mapping Controls to Architectural Implementations

Controls are reusable across the organization and typically are matched to a framework, as discussed in sections 2.6.4.and 2.6.5. However, cross-mapping efforts do not just involve frameworks. Extending controls to a reference architecture[91] creates better visibility into the specific organizational and technical structures. When presenting findings to nonauditors, architectural depictions are more effective than an auditor's spreadsheet for showing where a control will apply. Several references exist, which are typically based on an organizational focus or a technology.

Figure 2.12—Controls and Compliance Document Relationships

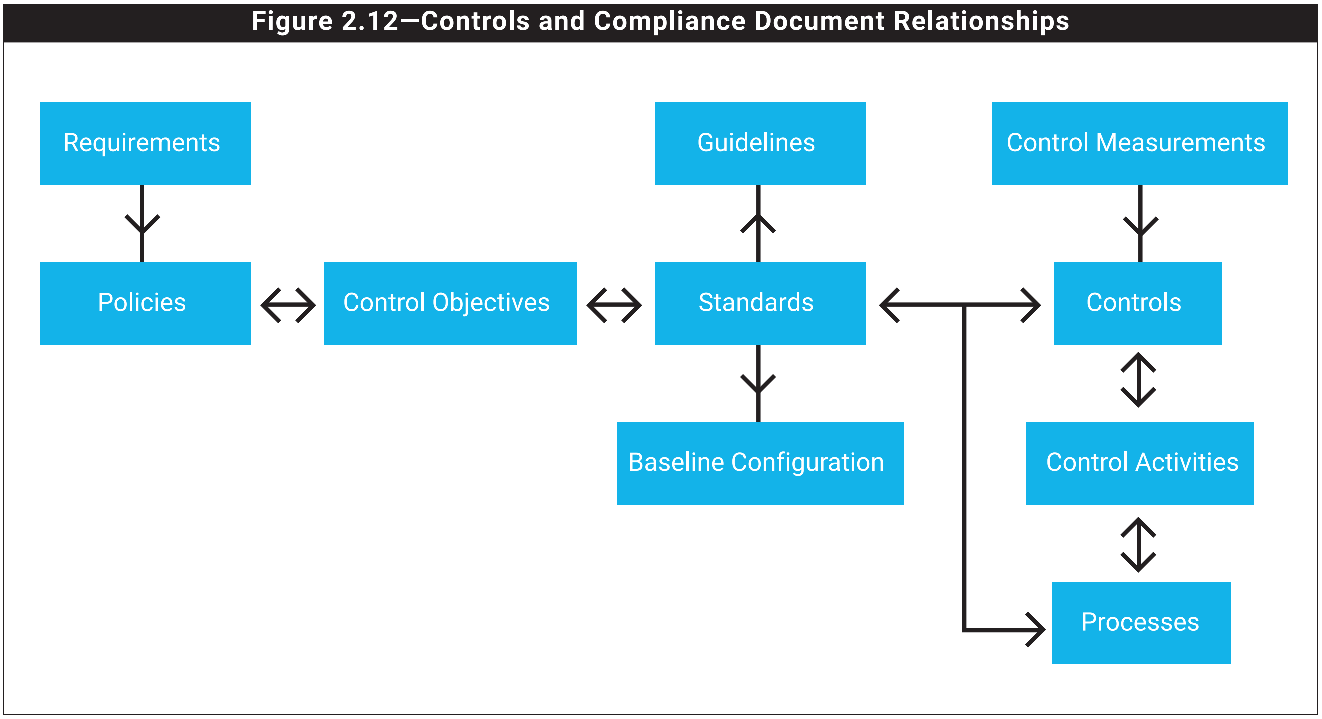

Defining Controls from Policies and Objectives to Activities and Controls

Before mapping architectures, it is essential to understand the relationships of control components (objectives, activities and controls) and associated governance documentation, identifying the control components from the top-level, strategic policy document to the implementation level processes and control measurements. **Figure 2.4** depicts the derivation of document and control components. By maintaining traceability through this chain, auditors may prove that implemented controls map to the starting requirements. The relationships within **figure 2.12** follow:

- Requirements (section 2.5) define Policies (sections 2.3.3.3 and 2.4.1.3) and influence control objectives.
- Control objectives (section 2.6.3.1) map to policies.
- Organizational standards map to control objectives.
- Baseline configurations flow from organizational standards and define platform-specific functionality.

[91] Reed, P.; "Reference Architecture: The best of best practices," IBM, 15 September 2002, www.ibm.com/developerworks/rational/library/2774.html

- Control activities (section 2.6.3.1) are actions (generally described in policies, standards and procedures) that help management mitigate risk to ensure the achievement of objectives.[92]
- Controls map to a corporate standard.
- Controls Measurements (section 2.8) map to controls.

Many-to-Many Mappings

Requirements introduced through policy, designed into baselines and eventually measured through metrics should be tracked for performance. Typically, as part of systems engineering, a many-to-many mapping exists, with specific controls satisfying multiple requirements. The most direct method of tracking comes in the form of a requirements traceability matrix (RTM). Although RTMs are common within government contracting, they are not found as frequently in enterprise environments. An RTM may include feature requests from current customers and necessary implementation requirements. Features promised to customers through contracts may be tracked with requirement IDs (**figure 2.13**). The test cases may involve interplay—e.g., to have feature A, feature B is a prerequisite—leading to additional regression analysis.

Figure 2.13—Sample Traceability Matrix

Req IDs	Req Tested	CCM AIS-04	CCM BCR-11	CCM EKM-04	CCM HRS-09	CCM IAM-01	PCI DSS 1.5	PSI DSS 2.2
Test Cases	18							
1.1.1	2	X		X				
1.1.2	1				X			
1.2.1	3		X			X		X
1.2.2	2	X					X	
1.2.3	2		X		X			
2.1	1					X		
2.2	4			X	X	X		X
3.1.1	1						X	
etc.	...							

Although RTMs are most easily understood as a method of guaranteeing coverage, software beyond the spreadsheet allows faster compliance audits and test case coverage. These software products may also be used past requirements tracking and focus on the broader GRC program.

Architecture and the Cloud

This section describes three types of cloud architectures.

Organizationally Focused Architectures

The US federal government adopted the NIST original Enterprise Architecture Model[93] in the 1980s. The Generalized Enterprise Reference Architecture and Methodology (GERAM)[94] of the 1990s extended enterprise

[92] Schandl, A.; P. Foster; "COSO Internal Control – Integrated Framework: An Implementation Guide for the Healthcare Provider Industry," Crowe, January 2019, www.coso.org/Documents/COSO-CROWE-COSO-Internal-Control-Integrated-Framework.pdf

[93] National Institute of Standards and Technology, "PM-7 ENTERPRISE ARCHITECTURE," National Vulnerability Database, https://nvd.nist.gov/800-53/Rev4/control/PM-7

[94] NIST, "An Overview of GERAM," International Conference on Enterprise Integration Modelling Technology 1997, http://www.mel.nist.gov/workshop/iceimt97/ice-gera.htm

modeling and provided building blocks for engineering efforts. Business and IT collaborators at ISACA created COBIT®[95] (Control Objectives for Information Technology), which is intended for end-to-end business and IT functional integrations, and is now COBIT 2019. One of the longer-running information technology architectures, the Information Technology Infrastructure Library (ITIL),[96] began with efforts by the UK government to describe the IT landscape. Axelos, a public-private venture, owns the current (fourth) ITIL incarnation.

Technology Architectures

In the computer environment, architecture is the conceptual structure and logical organization of a computer or computer-based system. The data center layout diagrams describing physical devices and their connections fit this description. There are multiple architectures within the IT space. Administrators reference abstractions from physical devices to component groups or block network designs as architectures. Software developers may be focused on an individual technology, such as the Enterprise Edition of the Java[97] product. All these examples point to constrained organizational units or limited environments.

Cloud Architectures

CSPs trend in the same direction, often describing their solution components as a complete architecture solving all use cases. AWS created the Well-Architected Framework[98] as a guide for implementing its service catalog with a focus on cost, reliability, security and performance. Microsoft released a Cloud Adoption Framework to "create and implement the business and technology strategies necessary for your organization to succeed in the cloud."[99] The Google Cloud Platform™ directs users to the best designs in GCP migration.[100]

CSA Trusted Cloud Initiative

Similarly, the CSA Trusted Cloud Initiative sought to create a holistic architectural model for a complete set of necessary cloud components. The project addressed not just the confined functionality of ITIL or cybersecurity, but the broad interpretation of business, IT, security and technology services. The project continued evolving, eventually becoming the CSA Enterprise Architecture (EA)(**figure 2.14**), with a broader understanding of the ubiquity of cloud. NIST adopted the EA in NIST SP 500-299 as part of its cloud reference architecture, solidifying the importance of the CSA approach.

EA Components

The Enterprise Architecture (EA) consolidates and coordinates features from four organizational architectures (**figure 2.14**).

[95] ISACA, "COBIT: An ISACA Framework," https://www.isaca.org/resources/cobit
[96] Axelos, "ITIL – IT service management," www.axelos.com/best-practice-solutions/itil
[97] Oracle, "Java EE at a Glance," www.oracle.com/technetwork/java/javaee/overview/index.html
[98] AWS, "AWS Well-Architected," https://aws.amazon.com/architecture/well-architected/
[99] Microsoft Azure, "Microsoft Cloud Adoption Framework for Azure," https://azure.microsoft.com/en-us/cloud-adoption-framework/
[100] Google Cloud, "Migrate for Compute Engine architecture," https://cloud.google.com/migrate/compute-engine/docs/4.10/concepts/architecture/gcp-reference-architecture

Figure 2.14—CSA Enterprise Architecture 2.0

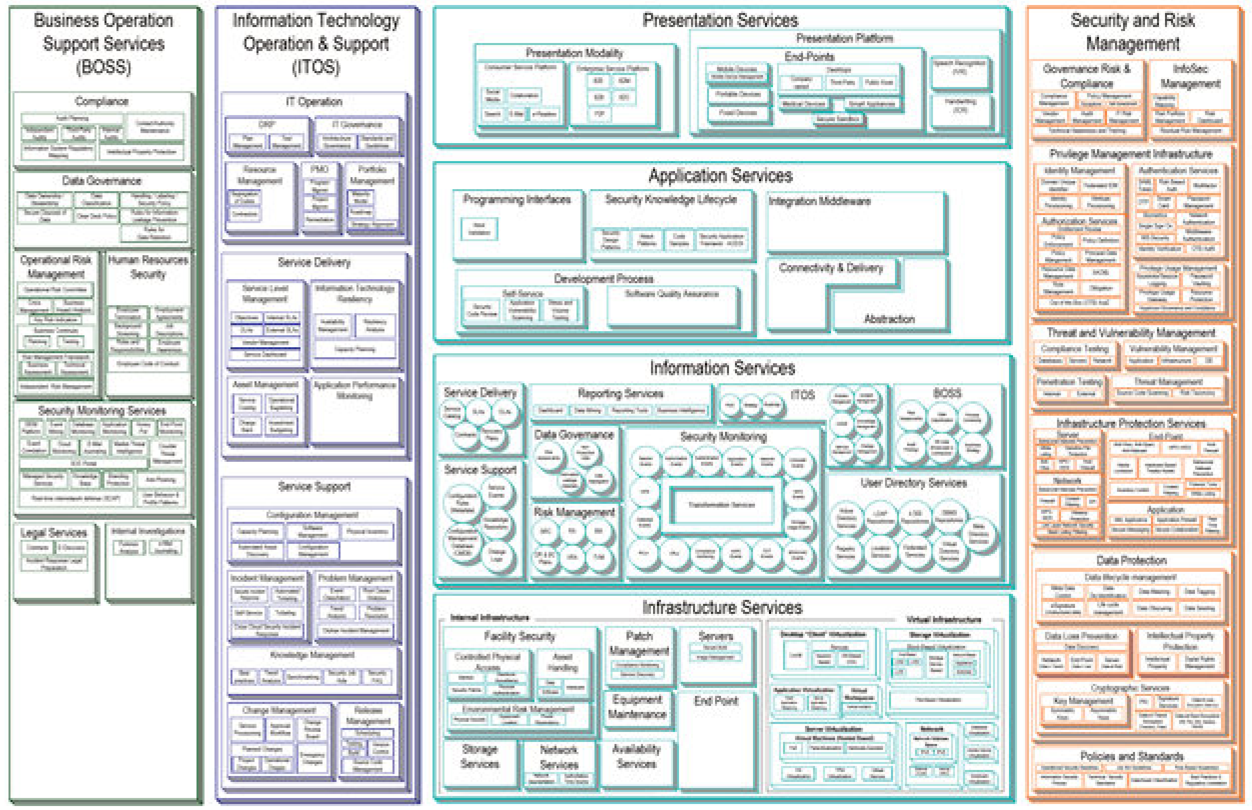

Source: Cloud Security Alliance, Enterprise Architecture 2.0

The EA guiding principles enable users to develop cross-platform capabilities and patterns, define protections that enable trust in the cloud, and provide direction to secure regulated information, all with an expectation of elastic, flexible and resilient support for multitenant and multilandlord platforms. The EA four domains and its ideas/best practices/models genesis follow:

- Business Operation Support Services (BOSS)—Sherwood Applied Business Security Architecture (SABSA)[101]
- Information Technology Operation Services (ITOS)—Information Technology Infrastructure Library (ITIL)[102]
- Technology services, including infrastructure (InfraSrv), information (InfoSrv), application (AS) and presentation (PS) services—The Open Group Application Framework (TOGAF)[103]
- Security and Risk Management (SRM)—The Jericho Forum[104]

The Enterprise Architecture consolidates the common components of these models, relying on interactions with the other three components to establish a tighter relationship for the whole. The EA can provide a presentation layer for control coverage and a communication medium for auditors, engineers, and security and management professionals.

[101] SABSA Enterprise Security Architecture, https://sabsa.org
[102] *Op cit* Axelos
[103] The Open Group, "The TOGAF Standard, Version 9.2 Overview," www.opengroup.org/togaf
[104] The Open Group, "Jericho Forum® Identity Commandments Reference: W125," https://publications.opengroup.org/w125

Tying the EA to the CCM

The all-encompassing topology of the Enterprise Architecture provides a perspective on the cloud. Because controls are reusable, there are alignments between the Enterprise Architecture and frameworks, such as NIST, ISO or the Cloud Controls Matrix. As parts of the CSA toolkit, the EA working group went with the CCM framework. The mapping constitutes an attempt to quantify the relationships of the 359 components in the EA and the 133 controls in the CCM: a 48K cell CCM to EA. There are perspectives and edge cases within the CCM to EA mapping spreadsheet[105] to understand, which necessitates a usage guide.[106] The mapping should not be used as gospel, as CCM entries may provide only partial detective, corrective or preventive controls over a particular EA component. The mapping offers users details on where to look, but does not document the nuances of completeness of control for that component. Where architectural elements show no coverage, organizations should look for gaps in control objectives.

Gap Identification Methodologies

A comparison of any two architectural designs or framework perspectives reveals differences. Gaps occur, even in holistic frameworks, complete architectures and extensive regulatory requirements. There is a natural inclination to immediately remediate a finding with another arbitrary control. Instead, organizations may progress from the policy level with purpose, and work down to final controls and measurements. **Figure 2.12** depicts the intersections of relationships between policies, standards, guidelines and processes with control objectives, control activities and control implementations. One example of a relationship between the Application and Interface Security domains is the Data Security and Integrity control group (AIS-04). The Enterprise Architecture containers and processes show 125 mapping suggestions for review in the EA to CCM mapping. See the example in **figure 2.15**.

Figure 2.15—Enterprise Architecture GRC Example

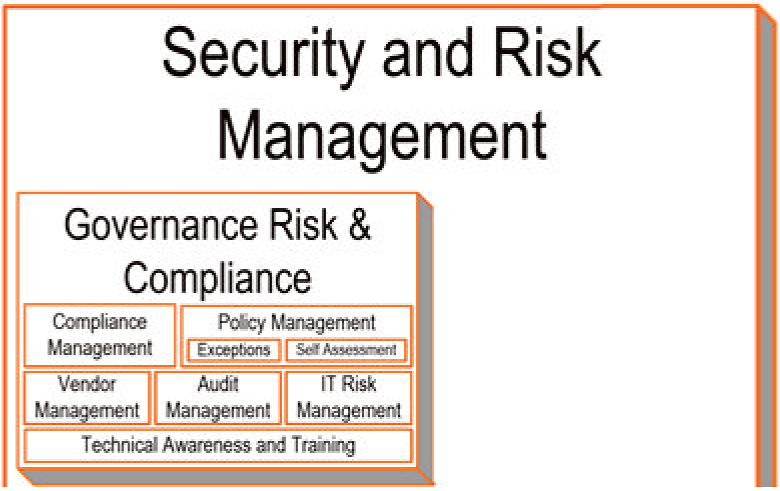

Source: Cloud Security Alliance, Enterprise Architecture 2.0

Working down from the policy level in the CCM, the Application and Interface domain mentions the following:

- AIS-04: Application and Interface Security → Data Security/Integrity
  - "Policies and procedures shall be established and maintained in support of data security…"[107]

[105] CCM to EA Mapping
[106] CCM to EA Mapping Usage Guide
[107] CSA Enterprise Architecture Reference Guide 2020

The AIS-04 row 11 maps to 125 items covering aspects of the enterprise architecture, e.g.:

- Security and Risk Management → Governance Risk and Compliance → Compliance Management
  - "Analyzes compliance with all specified internal information security policies, control standards and procedures."[108]

There are several policy callouts throughout the frameworks, CCM and architectures. The CCM AIS-04 and EA Security and Risk Management Requirement examples both call for policies. All policies fulfill requirements, although none of the tools listed identify what should serve as descriptive materials (discussed further in section 2.7.3). Requirements must be pulled from a variety of organizationally specific laws and contracts, (**figure 2.12**, section 2.6.6). Then, control objectives can be set for this policy and mapped into company standards. Those standards create a baseline configuration for the environment. The EA can help with the baseline configuration, visually identifying missing elements.

Figure 2.16—Example of a Tailored Enterprise Architecture With Coverage Gaps

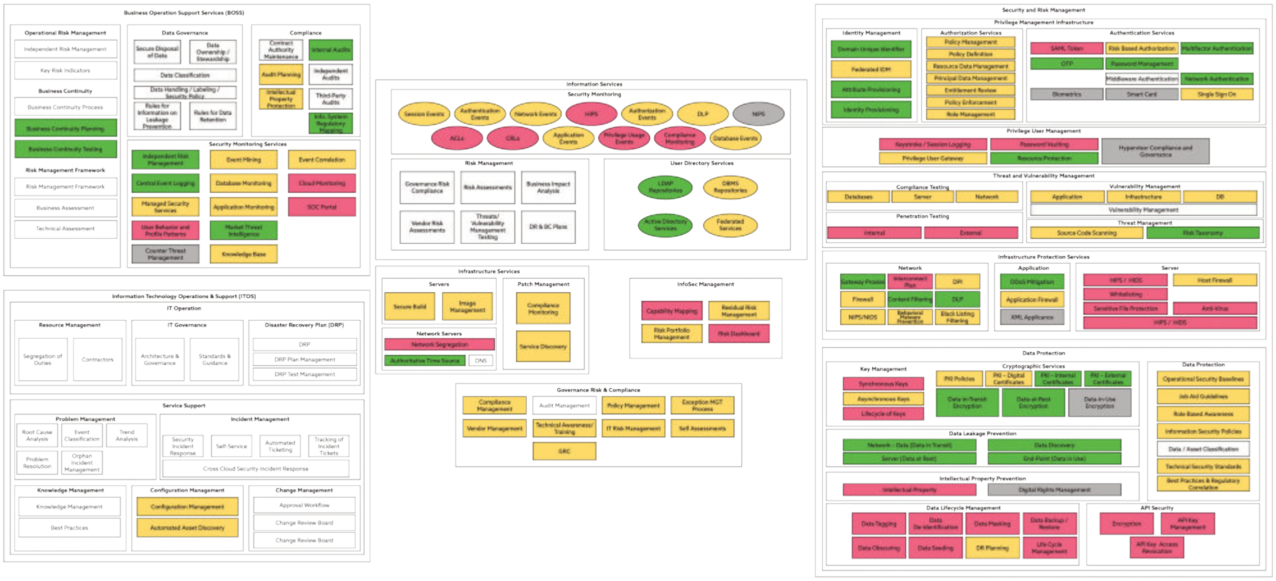

Source: Cloud Security Alliance, Enterprise Architecture 2.0

In **figure 2.16**, example of a tailored EA, color-coded progress indicators identify gaps within the four domains. As a presentation medium, this represents to the security professional that several deficiencies exist. Drilling down with the organizational owners furthers understanding. In the expanded detail mapping shown in **figure 2.17**, Legal Services and Internal Investigation represent reasonably covered architectural areas.

[108] *Ibid.*

Figure 2.17—Element Mapping: Business Operations Support Services (BOSS)

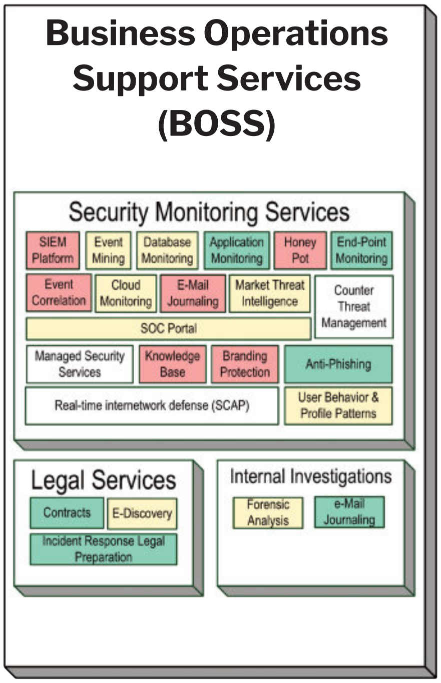

Source: Cloud Security Alliance, Enterprise Architecture 2.0

Several coverage gaps and tailoring activities appear within the Security Monitoring Services container, which can prompt further discussions with the organizational teams (security operations, legal and incident response).

2.8 Identifying Technical and Process Controls

2.8.1 Operationalizing Security Controls

The security controls found in standard control frameworks, like ISO/IEC 27001 or the CSA CCM, describe security measures or objectives in a technology-neutral language that is universally applicable to most information systems. For implementation in a real-life information system, these controls need to be translated into technical, organizational and physical measures tailored to the underlying platform, architecture, organizational structure and governance of the target information system.

The translation of standard security controls may take many forms:

- Updating procedures and their supporting documentation to address the organizational dimension of controls for business processes and human resources
- Reviewing the organization governance to ensure that management is in command of information security
- Reviewing contracts and adding appropriate clauses in contracts with entities in the supply chain and employees to cover identified risk
- Selecting security tools, technologies and relevant vendors to support control requirements for the confidentiality, integrity and availability of information system assets

- Addressing physical and environmental risk with appropriate technologies, processes and personnel
- Training and awareness

The operationalization of controls must be scrutinized. It can create security gaps or a false sense of security if the concrete security controls do not cover the higher-level security requirements that they are intended to embody.

2.8.2 Cloud-Specific Controls

It is necessary to translate generic, technology- and vendor-agnostic security controls into measures that are tailored to a specific technology, platform and architecture. The control operationalization process in the cloud may differ from traditional IT mainly for three reasons:

- The impact of the shared responsibility model ties the operationalization of a control to a combined effort from multiple parties (for instance, the CSP might need to expose an API for the customer to implement a certain monitoring control).
- Technologies have inherent differences (for example, while firewalls were a suitable solution for network segmentation and isolation in traditional IT, software-defined networking (SDN), workload isolation and security groups offer much more powerful alternatives in IaaS deployments).
- The scale and pace of change of the cloud service imposes the need to implement automated controls.

In general, it is important to avoid the temptation to replicate traditional IT practices in the cloud when better cloud-native approaches exist.

Controls from a Cloud Perspective

In traditional IT environments, emphasis is placed on people executing processes to fulfill corporate governance requirements. Authentication and authorization are assigned to roles with the familiar titles of administrator" or user. Email requests and managerial authorizations verify needs and budgets. Perimeter protections, such as firewalls and data loss prevention (DLP), contain specific threats to network connectivity, including administrator activity.

In contrast, the main motivators of CSPs include automation, speed and, to some extent, portability. Service-specific accounts (i.e., object storage, cryptography and compute) and attribute-based access control (ABAC) provide least-privilege methods to counter threats that do not exist within the enterprise. The velocity-of-execution goal precludes man-in-the-loop processes, avoiding delays through scripting, notifications and message queues. Whereas firewalls prevent network probes, decoupled services in cloud-native designs necessitate corrective controls. These controls are typically reactive to stimuli without human interaction.

The contrast between traditional IT and cloud environments exhibits itself most readily in the cloud metastructure. The cloud metastructure, as defined in CSA Guidance V4, includes “the protocols and mechanisms that provide the interface between the infrastructure layer and the other layers. The glue that ties the technologies and enables management and configuration.”[109] It includes the management plane components, which are network-enabled and remotely accessible. This waterline in the shared responsibility model defines where the CSP duties end and the consumer’s begin. It varies based on the service, platform, infrastructure (SPI) model, with the IaaS services most resembling typical enterprise controls. As organizations relinquish more operational aspects to the CSP, the technical and process controls become fewer, but harder to manage.

Metastructure aspects are vendor-specific, and necessitated controls cannot be readily moved for existing enterprise integration. Asking a public CSP to trust an enterprise active directory or LDAP controls creates a situation ripe for fraud. A CSP identity and access management (IAM) components require flexibility from the smallest of individual

[109] Cloud Security Alliance, “Security Guidance v4.0,” 26 July 2017, https://cloudsecurityalliance.org/research/guidance/

projects to the largest financial services and commercial undertakings. Organizations must adapt to new cloud-specific norms. An example of control adaptation is the proliferation of single sign-on (SSO) options or universal 2nd factor (U2F) credential storage integration devices.

In the cloud, controls use new capabilities. For example, the introduction of new categories of services, such as Function as a Service (FaaS),[110] enhances the concept of scripting. This allows automatic creation of new controls, a capability that is not necessarily available in an on-premises infrastructure. For instance, a FaaS corrective control may replace deleted credentials after notification; this does not have an enterprise analog. A preventive FaaS infrastructure control can proactively adjust security group administrative port assignments (i.e., SSH or RDP) from world open (0.0.0.0/0) to the IP range constraints of an organization.

Mitigating Cloud Controls

Not all cloud control offerings meet organizational needs. While operationalizing a control, the cloud customer might face different possible options that are combined in different control implementation scenarios, as follows:

- The CSP can offer security capabilities that fully satisfy the customer requirements.
- The CSP can meet most of the adoption requirements, although there may be unacceptable feature implementations that require additional mitigations. Several organizations within the financial services industry use the cloud but are not comfortable with the cryptography solutions that the CSP offers. Whether they better trust their overall cryptosystem or because they are forced to directly manage the encryption keys for compliance reasons, these customers might decide to look at enterprise implementations such as onsite hardware security modules (HSMs).
- The CSP does not integrate with other solutions and requires specific third-party workarounds or custom application development.
- The CSP provides a rich set of APIs that can be used by the cloud customer or third-party security vendor (e.g., CASB) to develop security capabilities in addition to the ones offered by the native cloud service (e.g., security monitoring solutions and DLP).

Cloud Deployment Model Impact on Control Operationalization

The operationalization of controls may vary considerably between different deployment models, especially for technical controls. For example, in the CSA Cloud Controls Matrix (see chapter 3) and the CCM Shared Responsibility Mapping,[111] it is possible to identify a variation between CSP, shared and consumer. On a pure IaaS platform, 100 of the controls remain under the responsibility of the cloud customer, while 129 are the CSP responsibility. This means that the customer has much flexibility in the operationalization of controls and can select among different solutions to best match its requirements.

On a SaaS platform, most of the responsibility is with the CSP, with 130 controls as its responsibility, and only 83 controls are directly controlled by the consumer. In these cases, the CSP falls under the supplier relationship controls and the SaaS provider has responsibility to confirm that it has adequately validated that the CSP's controls meet its required specifications. The operationalization of controls is usually constrained by the security tools and configuration options embedded in the platform. The downside is that these tools and configuration options may not

[110] See IBM Cloud Education, "What is FaaS (Function-as-a-Service)?," 30 July 2019, www.ibm.com/cloud/learn/faas: "FaaS (Function-as-a-Service) is a type of cloud-computing service that allows you to execute code in response to events without the complex infrastructure typically associated with building and launching microservices applications. Hosting a software application on the internet typically requires provisioning and managing a virtual or physical server and managing an operating system and web server hosting processes. With FaaS, the physical hardware, virtual machine operating system, and web server software management are all handled automatically by your cloud service provider. This allows you to focus solely on individual functions in your application code."

[111] Google Sheets, https://docs.google.com/spreadsheets/d/1KErVfDWx3JQL4QlHYE5CBSYhMe0Le_K52FprvJyr5uY/edit

match the requirements of the specified control, and there may need to be additional compensating controls in some cases.

2.8.3 Evaluating Specific Technical Controls

Identifying Sources for Controls

The selection of an authoritative reference source for creating a baseline configuration depends on where the organization is in the cloud migration process. If a consumer has not chosen a CSP and is in the evaluation phase, industry organizations can help narrow the field through benchmark control lists and features comparisons. If the CSP decision has already been made, vendor requirements and vendor-specific product benchmarks may assist in creating a secure implementation.

Industry Organizations

In addition to vendors, industry groups with specific focuses may provide direction. The SANS Institute put together a large library of early compliance controls literature. Although the institute still produces a large volume of white papers, it handed over the Top 20 Critical Controls to the Center for Internet Security (CIS)[112] for upkeep. CIS began generating new collaborative benchmarks, working with CSPs and other groups, such as the CSA, for product- and technology-specific implementation methodologies. Application developers pay attention to the Open Web Application Security Project (OWASP) and its Top 10.[113]

Center for Internet Security® (CIS®)

The Center for Internet Security offers security controls and cloud-relevant benchmarking. The benchmarking provides prescriptive guidance for establishing a secure configuration posture on many technology- and vendor-specific services and products, including the following:

- Cloud Providers:
  - Amazon Web Services® (AWS)
  - Microsoft® Azure
  - Google Cloud Platform™
- Virtualization:
  - Docker®
  - Kubernetes®
  - VMware® server

The CIS benchmarking includes several other categories:

- Database servers
- Desktop servers
- Web servers
- Networking devices
- Operating systems

[112] Center for Internet Security, "The 20 CIS Controls & Resources," www.cisecurity.org/controls/cis-controls-list/
[113] OWASP, "OWASP Top Ten," https://owasp.org/www-project-top-ten/

Open Web Application Security Project (OWASP)

"The Open Web Application Security Project® (OWASP®) is a nonprofit foundation that works to improve the security of software."[114] Its community-led software projects are open source. One of its best-known projects, the OWASP Top Ten, is a standard awareness guideline that is "globally recognized by developers as the first step towards more secure coding."[115] Many cloud services incorporate the Top Ten into their product offerings, including major web application firewalls (WAFs): AWS WAF, Microsoft Azure Application Gateway WAF and Google Cloud Armor WAF.

CSA STAR

STAR is the CSA security and privacy governance and compliance program. Consumers may use it for understanding CSP-implemented controls through CCM or Consensus Assessments Initiative Questionnaire (CAIQ) evaluations and for identifying shared responsibilities. See chapter 9 for further details.

Vendor Sources

The best sources of information for IaaS and PaaS services are the CSP documentation repositories. Mature CSPs include shared responsibility mappings against common frameworks, such as the AWS Standardized Architecture for NIST-based Assurance Frameworks.[116] As a product consumer, CSPs should share under NDA their SOC II (Service Organization Controls II) type 2 assessment findings. A SOC II example[117] from the Microsoft library requires customer sign-in prior to download. These documents should direct users to customer control responsibilities. Additionally, it is customary to have CSP experts onsite to assist further with implementing and configuring these controls.

Determining Baseline Configurations

The standards and associated platform-specific controls required for baseline configurations are best defined in organization security policies. Security policies are abstract, directing security control objectives and not specific products or implementations. For example, a security policy may state that the organization must protect data at rest with encryption. In response, as part of its corporate standard, the organization may require virtual machine hard disk encryption. The option of following the standard by implementing encryption on its own may be technically infeasible or unreasonably expensive. Instead, the company may consider using CSP-managed services for encryption controls. Each CSP has its own dedicated service that aims to satisfy customer requirements and control objectives.

CSP Product Comparisons

Depending on the products and CSP, similarities and differences abound. Although there are many generic cloud concepts, the implementation decisions determine the details. CSPs decide on features based on business decisions pursuing vertical sectors, satisfying customer requests, continuing legacy support, or meeting legislative requirements. Public CSP offerings typically satisfy most of the demand, with certifications for healthcare, financial services or energy regulation (see section 2.6). However, there are special community clouds set up with default controls that differ significantly from public options, for example:

[114] OWASP, "Who is the OWASP® Foundation?," https://owasp.org
[115] OWASP, "OWASP Top Ten," https://owasp.org/www-project-top-ten/
[116] AWS, "Standardized Architecture for NIST-based Assurance Frameworks on the AWS Cloud: Quick Start Reference Deployment," January 2016, https://docs.aws.amazon.com/quickstart/latest/compliance-nist/welcome.html
[117] Microsoft, "Service Organization Controls (SOC)," 22 September 2020, https://docs.microsoft.com/en-us/microsoft-365/compliance/offering-soc?view=o365-worldwide

- **AWS GovCloud**[118]—Includes International Traffic in Arms Restrictions (ITAR)[119] requirements
- **Microsoft G-cloud**[120]—United Kingdom government cloud includes GDPR restrictions
- **AWS Intelligence Cloud**[121]—US-based intelligence agencies meet special hardware implementations and connectivity requirements
- **Microsoft Joint Enterprise Defense Infrastructure (JEDI)**[122]—US Department of Defense meets special hardware and connectivity requirements
- **Google Cloud for Government**[123]—Offers worldwide government agency support
- **Other CSPs**—Target specific community spaces with wholly focused products, e.g., the Armor[124] platform

A complete analysis of all the CSP options is beyond the scope of this study guide.

Examples of product implementations and the design requirements that they satisfy are shown in the following application protection and encryption requirements analysis.

Application Protection Requirements Analysis and Baseline Configuration

This is a common requirement set for mitigating software vulnerabilities and external network threats to an application. Organizations can implement their own instance-based web application firewall (WAF) or outsource it to a third-party CASB. The complexity, cost and additional supply chain risk may not prove of benefit to the organization when compared with the CSP service implementation.

Sample application access protection requirements include the following:

- A WAF shall be used to protect data with a critical data classification rating
- A rule set shall be used based on OWASP Top Ten[125]
- Additional WAF rules may be added when deemed necessary
- WAF shall utilize alerting with company notifications monitoring
- A shared responsibility model shall be in place

Azure

The application access protection requirements for Azure include implementing Azure WAF on Application Gateway, with the standard rule set available.[126]

[118] AWS, "Introduction to the AWS GovCloud (US) Regions," https://aws.amazon.com/govcloud-us/

[119] US Department of State, "The International Traffic in Arms Regulations (ITAR)," Directorate of Defense Trade Controls, www.pmddtc.state.gov/ddtc_public?id=ddtc_kb_article_page&sys_id=%2024d528fddbfc930044f9ff621f961987

[120] Microsoft, "United Kingdom Government-Cloud (G-Cloud)," https://docs.microsoft.com/en-us/microsoft-365/compliance/offering-g-cloud-uk

[121] AWS, "Cloud Computing for the U.S. Intelligence Community," https://aws.amazon.com/federal/us-intelligence-community/

[122] Inspector General, "Report on Joint Enterprise Defense Infrastructure (JEDI) Cloud Procurement," US Department of Defense, 13 April 2020, https://media.defense.gov/2020/Apr/21/2002285087/-1/-1/1/REPORT%20ON%20THE%20JOINT%20ENTERPRISE%20DEFENSE%20INFRASTRUCTURE%20(JEDI)%20CLOUD%20PROCUREMENT%20DODIG-2020-079.PDF

[123] Google Cloud, "Google Cloud for government," https://cloud.google.com/solutions/government

[124] Google Cloud, "Google Cloud Armor," https://cloud.google.com/armor

[125] *Op cit* OWASP

[126] Microsoft, "What is Azure Web Application Firewall on Azure Application Gateway?," 16 September 2020, https://docs.microsoft.com/en-us/azure/web-application-firewall/ag/ag-overview

AWS

The application access protection requirements for AWS include implementing AWS WAF, with the standard rule set available.[127]

In this example, both CSPs maintain a similar product. This can principally be due to well-established enterprise products that translate effectively to the cloud. There are similar threat models and use cases identified within both providers.

Encryption Requirements Analysis and Baseline Configuration

This is a requirement commonly found in policies and internal control frameworks based on legislative and international standard requirements. The purpose is to address confidentiality and availability aspects within their control implementations, e.g., by extending the virtual disk encryption baseline configurations.

Sample encryption requirements follow:

- All data encrypted at rest for all storage services
- Keys managed in a vault
- Vault backed by an HSM
- Keys rotated every 90 days
- Shared responsibility model in place

Azure

The encryption requirements for Azure include the following:

- Implement Azure Disk Encryption.[128] This service is integrated with Azure KeyVault.[129]
- Use Azure Security Center to monitor for alerts on VMs that are not encrypted.

AWS

The encryption requirements for AWS include the following:

- Implement encryption on the storage volumes used by the EC2 server.
- Amazon EBS encryption is an encryption solution for EBS volumes and snapshots.[130] It uses AWS Key Management Service (AWS KMS) customer master keys (CMKs).[131]
- Activate an AWS Config[132] rule to check for Not Encrypted EC2 Volumes.

In the AWS example, CSP implementation differences require additional set-up and per-use costs. Consumer business models and common business cases can influence product choices. Ease of use and system flexibility may push a consumer toward fewer service needs in Azure, as would legacy compliance for an all-Microsoft enterprise structure.

[127] AWS, "AWS WAF - Web Application Firewall," https://aws.amazon.com/waf/
[128] Microsoft, "Azure Disk Encryption for virtual machines and virtual machine scale sets," 15 October 2019, https://docs.microsoft.com/en-us/azure/security/fundamentals/azure-disk-encryption-vms-vmss
[129] Microsoft Azure, "Key Vault," https://azure.microsoft.com/en-us/services/key-vault/
[130] AWS, "Amazon EBS encryption," https://docs.aws.amazon.com/AWSEC2/latest/UserGuide/EBSEncryption.html
[131] AWS, "Customer master keys (CMKs)," https://docs.aws.amazon.com/kms/latest/developerguide/concepts.html#master_keys
[132] AWS, "AWS Config," https://aws.amazon.com/config/

Additional Controls Considerations

Beyond the previous simplified requirements analysis example, specific controls implementations vary based on cloud delivery model, specific technology or vendor construction methods. This section includes some things to examine.

Delivery Model Considerations

Control implementations for cloud delivery model follow:

- **IaaS**
    - **Consider**: Multifactor authentication (MFA) for an account may require a specific token vendor implementation that is incompatible or requires manual translation for use with existing cryptographic management systems.
    - **Cloud effect**: Complicated policy, standards and processes may insist on enterprise key controls based on a competing technology.
        - **AWS**—SafeNet Gemalto tokens[133]
        - **Google**—Titan Security Keys[134]
        - **Microsoft**—OATH TOTP hardware tokens[135]
- **PaaS**
    - **Consider**: PaaS database feature implementations vary from service to service. The attractiveness of PaaS is based on increasing resiliency and speed, or on automating difficult activities. For example, master/slave node elections may be necessary for clusters in an on-premises or IaaS deployment. PaaS services may build in clustering features, removing the need for election configuration.
    - **Cloud effect**: Use of these configurations may introduce stickiness, the idea that scripting code and feature automation may not be portable and may result in vendor lock-in. Deciding a direction should be a strategic decision, weighing ease of use versus provider dependence. Following are three references for Oracle as a Service:
        - Oracle Database as a Service[136]
        - Amazon RDS for Oracle[137]
        - Oracle solutions on Microsoft[138]
- **SaaS**
    - **Consider**: Data loss prevention controls within enterprise storage environments utilize techniques for data at rest, data in motion and data in use. These include capabilities of exact file matching, partial indexed data matching and pattern matching. Within a cloud environment, and particularly with SaaS providers, the shared responsibility model limits data access. Storage vendors may only provide a subset of pattern matching, such as 16-digit credit card number checks, 9-digit US Social Security Numbers, or filename/list checks.

[133] AWS, “Multi-factor Authentication,” https://aws.amazon.com/iam/features/mfa/

[134] Google Cloud, “Titan Security Key,” https://cloud.google.com/titan-security-key

[135] Microsoft, “What authentication and verification methods are available in Azure Active Directory?,” https://docs.microsoft.com/en-us/azure/active-directory/authentication/concept-authentication-methods

[136] Oracle, Oracle Enterprise Manager Cloud Administration Guide,” https://docs.oracle.com/html/E28814_01/cloud_db_overview.htm

[137] AWS, “Amazon RDS for Oracle,” https://aws.amazon.com/rds/oracle/

[138] Microsoft, “Overview of Oracle Applications and solutions on Azure,” https://docs.microsoft.com/en-us/azure/virtual-machines/workloads/oracle/oracle-overview

- **Cloud effect**: If previously implemented controls or processes required free-form or index matching, users may need to consider additional CSP services, instance-backed products, CASBs, or other third-party integrations. Several storage platform or Storage as a Service combinations:
    - AWS S3 and Macie[139]
    - Google Cloud Storage and Cloud DLP[140]
    - Microsoft Azure Blob and Information Protection[141]
    - Box DLP[142]
    - Dropbox with Symantec DLP[143]
    - Zscaler Cloud DLP[144]
    - Netskope DLP[145]

CSP Feature Considerations

The previous baseline configuration examples show a simplified requirements analysis, but control implementation choices may not be as straightforward within a single CSP. The volume of cryptographic choices per service varies. Examining key creation and management within IaaS environments should include the following:

- **Consider**: Enterprise solutions require a key management server and client-based access capabilities for encryption.
- **Cloud effect**: Consumers could continue to use their existing key management infrastructure or stand up an IaaS key management server instance, facing most of the same control challenges. Otherwise, three basic scenarios exist with cloud encryption, all surrounding who controls the key material and where the cryptographic operation occurs. Encryption in the cloud methodologies begin with who maintains the key material. General cloud encryption options:
    - **a.** Server-side encryption (SSE) and key management is the simplest format for implementation. The CSP handles everything. The overhead for doing anything beyond this through other services (like object storage) is unreasonable. Think of SSE as a check mark in a CSP management console to turn on encryption.
    - **b.** If the CSP offers a customer master key (customer managed key or bring your own key), the consumer uses its key management server for key generation. The customer loads the key into the CSP system and the CSP handles all encryption operations based on that value. The key could be exported or removed and replaced within the service. Vendor key management procedures must be reviewed for customer control acceptance.
    - **c.** A hardware security module (HSM) has a special designation for encryption functions.[146] It is the most secure option, and hardest for the CSP implementation and the customer's use. Once a key goes in, it cannot be exported—it can only be replaced (even if it's with the same key after removal). Per HSM design, any attempts to remove the key will erase the device.

Examples of multiple technology implementation names within a services vendor follow:

- AWS key management service (KMS):
    - Server-side encryption (SSE) SSE-KMS with simple storage service (S3)[147]

[139] AWS, "Amazon Macie," https://aws.amazon.com/macie/
[140] Google Cloud, "Inspecting storage and databases for sensitive data," https://cloud.google.com/dlp/docs/inspecting-storage
[141] Microsoft, "Azure Information Protection documentation," https://docs.microsoft.com/en-us/azure/information-protection/
[142] Box, "Data loss prevention," www.box.com/security/data-loss-prevention
[143] Dropbox, "Symantec," www.dropbox.com/app-integrations/symantec
[144] Zscaler, "Data loss prevention has become more difficult," www.zscaler.com/products/data-loss-prevention
[145] Netskope, "Data Loss Prevention," www.netskope.com/products/capabilities/data-protection
[146] National Institute of Standards and Technology (NIST), FIPS 140-3 *Security Requirements for Cryptographic Modules*, Computer Security Resource Center (CSRC), 22 March 2019, https://csrc.nist.gov/publications/detail/fips/140/3/final
[147] AWS, "How Amazon Simple Storage Service (Amazon S3) uses AWS KMS," https://docs.aws.amazon.com/kms/latest/developerguide/services-s3.html

  - KMS customer master key (CMK or bring your own key)[148]
  - Hardware security module (HSM) customer master key creation[149]
- Azure Key Vault:
  - Server-side encryption using service-managed keys[150]
  - Server-side encryption using customer-managed keys in Azure Key Vault[151]
  - Server-side encryption using customer-managed keys on customer-controlled hardware (host your own key [HYOK])[152]

As shown in the Delivery Model section (2.7.3.4.1), CSP implementations up the SPI stack remove options for users. Three additional examples from Oracle, Microsoft and CyberArk provide perspectives to reinforce the point that consumers may find another provider if the controls that a CSP chooses are unacceptable.

- Oracle as a Service employs transparent database encryption for its PaaS product. Server-side implementations use two-tiered keys for protecting the tablespace data, files and backups. Oracle avoids CMK or HSM implementations, opting for ease of use:
  - Oracle as a Service Transparent Database Encryption[153]
- While PaaS solutions do not require underlying IaaS building blocks, SharePoint as a Service uses a wide array of other Microsoft Azure services. Although that may be seen as depriving the consumer of choice, it provides opportunities for integration beyond the server-side encryption checkbox within the offering:
  - Microsoft 365 SharePoint as a Service[154]
- Cryptography as a Service is an example of encryption features and adjustments within a SaaS provider that go beyond a protect-my-data checkbox. Key storage, rotation and monitoring control options for SSE and CMK rival IaaS build-it-yourself implementations:
  - CyberArk Privilege Access Management (PAM) as a Service[155]

Technology Considerations

Although not an exhaustive services list, the following cloud technology topics support control creation in cloud terms.

- **Identity**
  - **Enterprise integration strategies**—Consider the cloud as always on, or at least that is the intent. Integration options that insist on using existing identity services will not have the same level of broad network access. Putting a complete or partial copy into the cloud, or a separate additional directory, may avoid service interruptions stemming from the enterprise. Cloud controls must be in place to keep the copies in-sync or the separate hosted directory available.
  - **Attributes and tagging**—Look past the identity components of the enterprise. The cloud offers attribute-based access controls (ABAC), providing additional activity confirmation for privileged actions, service accounts and data classification operations. Expect new control methods, tagging headaches, and the potential for complicated rulesets in environments such as data lakes.

[148] AWS, "Importing key material in AWS Key Management Service (AWS KMS)," https://docs.aws.amazon.com/kms/latest/developerguide/importing-keys.html
[149] AWS, "AWS CloudHSM," https://aws.amazon.com/cloudhsm/
[150] Microsoft, "Azure data encryption at rest," 13 August 2020, https://docs.microsoft.com/en-us/azure/security/fundamentals/encryption-atrest#server-side-encryption-using-service-managed-keys
[151] *Ibid.*
[152] *Ibid.*
[153] Oracle, "Administering Oracle Database Classic Cloud Service," https://docs.oracle.com/en/cloud/paas/database-dbaas-cloud/csdbi/data-security.html
[154] Microsoft, "Encryption for Skype for Business, OneDrive for Business, SharePoint Online, Microsoft Teams, and Exchange Online," 12 June 2020, https://docs.microsoft.com/en-us/compliance/assurance/assurance-encryption-for-microsoft-365-services
[155] Cyberark, "Cyberark Privilege Cloud," www.cyberark.com/products/privileged-account-security-solution/cyberark-privilege-cloud/

    - **Federation sources**—Multinational conglomerates can take a cue from the US Department of Defense in setting up their own central federated identity service, such as DEERS/RAPIDS.[156] Very few enterprises need many of the trust controls present for a military system, but the same root of trust idea occurs. Federation works well in the cloud through tokenization, but it will break down if the supply chain becomes unavailable and tokens expire.

- **Logging**
    - **CSP reliance**—Due to the shared responsibility model and metastructure waterline, a CSP management activity logging depends completely on the CSP capabilities and processes. The typical cloud automation mantra scripts actions in lieu of administrators. Logging thereby relies on provider maturity, confidence in maintaining multitenancy, preventing CSP intellectual property or trade secret data leakage, and general product goals.
    - **Forensic considerations**—Cloud logging depends on utility pricing models. CSPs do not benefit financially from over recording. Instead, they offer enhanced logs to the tenant at additional cost. Consumers should carefully consider what data are necessary for controls surrounding data retention compliance, incident response and chain of custody.
    - **Exiting the service**—The utility model comes into play again as CSPs discuss exporting logs. Large customers of Security Incident and Event Monitoring (SIEM) as a Service may find the network costs for exporting log files required for complete exit unjustifiable, and may accept the SIEM as another high-priced log repository.
- **Storage**
    - **Variable access and resiliency**—There are two variables in cloud storage. Moving bytes in the service frequently becomes expensive. Similar to moving to tape, having fewer accesses at a higher cost per access means less total cost. Secondly, guaranteeing that your data remains in the cloud requires more CSP equipment to account for the physical threats, mean time between failures (MTBF) and programming errors. Examples of possible adjustments:
        - AWS S3 Infrequent Access[157]
        - AWS Reduced Redundancy[158]
        - AWS Glacier[159]
    - **Destruction methods**—Data destruction requirements will change after cloud adoption. Without access below the waterline, degaussing or physical destruction are not valid options, and multipass wipes may prove inefficient or ineffective. Cryptoshredding may be the best course of action, depending on the deployment model. Consider meta tagging strategies for data types that may cause issues, specifically with respect to privacy and the right to be forgotten.
- **Containers**
    - **Source Provenance and Pedigree**—Containerization allows portability similar to automated configuration management without sacrificing the speed of vendor images. IaaS providers offer container flavors as a service, including AWS Docker,[160] AWS Elastic Kubernetes,[161] Azure Container Services,[162] Google Containers on Compute Engine[163] and Google Kubernetes Engine.[164] Typically set up as microservices, containers run others' code within a cloud customer's infrastructure. Controls should be placed in keeping with the sources and quality of the products used.

[156] HRC, "DEERS/RAPIDS/TASS Support Office," https://www.hrc.army.mil/content/DEERS~2FRAPIDS~2FTASS%20Support%20Office
[157] AWS, "Amazon S3 pricing," aws.amazon.com/s3/pricing/
[158] AWS, "Amazon S3 Reduced Redundancy Storage," aws.amazon.com/s3/reduced-redundancy/
[159] AWS, "Amazon S3 Glacier & S3 Glacier Deep Archive," aws.amazon.com/glacier/
[160] AWS, "What is Docker?," https://aws.amazon.com/docker/
[161] AWS, "Amazon Elastic Kubernetes Service," https://aws.amazon.com/eks/
[162] Microsoft Azure, "Container services," https://azure.microsoft.com/en-us/product-categories/containers/
[163] Google Cloud, "Deploying containers on VMs and MIGs," https://cloud.google.com/compute/docs/containers/deploying-containers
[164] Google Cloud, "Google Kubernetes Engine," https://cloud.google.com/kubernetes-engine

2.9 Measuring Effectiveness Through Metrics

2.9.1 Control Effectiveness

Information security management is an ongoing process that must continuously adapt to a changing threat landscape and evolving business objectives.

Monitoring, a critical component of information security, should encompass the following activities:

- Collect and review security events, preferably before they become security incidents.
- Monitor the effectiveness of information security processes and governance.
- Periodically review risk, because vulnerabilities, likelihoods, threats and even assets change.
- Monitor the effectiveness of implemented controls.

Monitoring the effectiveness of implemented controls means, in practice, evaluating the effectiveness of the concrete security controls that operationalize the organization control objectives.

Measuring the effectiveness of concrete security controls means that a monitoring process is in place to do the following:

- Perform continuous testing and collect security events, system logs and other evidence related to security.
- Perform measurements that take the collected data as input and produce qualitative or quantitative results to assess the effectiveness of each concrete security control.
- Establish objectives that contrast the measured results with expected results.

As a simple example, consider a technical control that requires users to pick stronger passwords by enforcing that they must be at least eight characters long and contain at least a number, a lowercase letter, and an uppercase letter. Whether the control is in place can be tested with an automated tool that regularly attempts to create an account with a password that does not match the requirements. This test should result in an error and rejection of the created account. Another option is to measure the number of monthly password security incidents reported to the support desk. The collected numbers should indicate whether the password policy reduces the number of password-related incidents, or if operationalization of the control is effective.

2.9.2 Metrics

> *If you cannot measure it, you cannot improve it.*[165]

The process of collecting evidence and assessing it to obtain a quantitative or qualitative result for the purpose of evaluating the performance of an information security system is essentially the process of using security metrics.

Formally speaking, as defined in ISO 19086-1, a metric is a standard for measurement that defines the rules for performing the measurement and understanding the results of a measurement. Note that security professionals sometimes use the word metric colloquially to describe measurement results, which can create confusion.

A measurement is defined as a process to quantify or qualify an attribute. In this context, a security attribute is a property or characteristic of an object that can be distinguished quantitatively or qualitatively by human or automated means, according to ISO/IEC 27000.

A measurement, as a process, involves the gathering of data, such as system logs, test results, configuration files, security events and sometimes the results of other measurements. These elements are often collectively referred to as

[165] Commonly attributed to Peter Drucker (1909-2005), a widely known and influential thinker on business management

evidence. ISO/IEC 27000 and many other sources refer to the result of a measurement as a measure. More recent initiatives, such as NIST SP 500-307, ISO/IEC 19086 and CSA STAR Continuous prefer the term measurement result, because the word measure has multiple meanings in information security and is a source of confusion when it comes to metrics. Measure and measurement result should be understood as equivalent.

Andrew Jaquith, widely regarded as one of the top security metric experts, defines[166] good metrics as having the following ideal characteristics:

- **Consistently measured**—Measurements should not involve subjective criteria.
- **Cheap to gather**—Gathering evidence and producing a measurement result should be as automated as possible.
- **Expressed as a cardinal number or percentage**—Avoid qualitative labels, such as high, medium or low.
- **Expressed using at least one unit of measure**—Examples include defects, hours and dollars.
- **Contextually specific**—Results are relevant, enabling decision makers to take action.

These are ideal characteristics of course, and practical metrics cannot always follow them, but they establish a good standard to target.

Recent research[167] suggests a few best practices:

- Avoid metrics expressing results as an average or a mean, which hides information.
- Prefer metrics expressing maximums or minimums.
- Prefer metrics expressing totals to percentages.

Consider a hypothetical metric describing the average time to respond to a severe incident. Assume there are 20 incidents, and 19 are treated within two hours, but one is treated in 15 days. The average time to process an incident is less than 20 hours. That result might not look too bad, but it hides the 15-day delay for one particular incident. Providing the maximum time to respond might be a better indicator.

A metric definition should[168] include at least the following:

- The security/privacy attributes to which it applies
- A definition of the evidence needed to assess the attribute
- A description of the possible ways relevant evidence may be collected
- A specification of the measurement method: a process that takes evidence as input and produces a measurement result as output
- The time span of the measurement (e.g., minutes, weeks, months)
- The required format and applicable units of the measurement result

There are many standards for controls in the cloud, but there are few standards or best practices related to security metrics,[169] especially[170] in the cloud.[171]

[166] Jaquith, A.; *Security Metrics: Replacing Fear, Uncertainty, and Doubt*, Addison-Wesley, 2007

[167] EU-SEC project, "Deliverable D1.4 Principles, Criteria and Requirements for a Multi-Party Recognition and Continuous Auditing Based Certifications," 3 June 2018, https://cdn0.scrvt.com/fokus/c56a737828fef8cf/79add0dec99a/D1.4-Multiparty-recognition-V1.0.pdf

[168] *Ibid.*

[169] NIST SP 800-55 is one exception.

[170] In 2020, to fill this gap, CSA launched an industry Working Group that aims to create a catalog of security metrics for the cloud, with mappings to the CSA CCM.

[171] NIST SP 500-307 is also an exception here.

2.9.3 Applying Metrics

Measurements performed according to well-defined metrics enable efficient and effective governance. Auditors provide value to stakeholders through the metrics and analysis provided. Results should be verifiable, based on quantitative metrics that are repeatable and that align with prioritized goals and objectives.

The collection of meaningful cloud measurement results within cloud service products is frequently accomplished through leveraging available automation and setting parameters.

Both the cloud customer and the auditors need to understand how measurement results are collected (automated or manually) and optimize the value of transforming data into information that can provide the basis for taking appropriate action.

Metrics-based decision making, whether by a human or through automation, requires a structured approach to ensure the right measurement results and metrics are used to gain meaningful and actionable insights. These insights provide transparency, helping ensure cloud services are adequately secured while satisfying user requirements.

A common way is to apply a Goal-Question-Metric approach (**figure 2.18**).

Figure 2.18—Goal-Question-Metric Approach

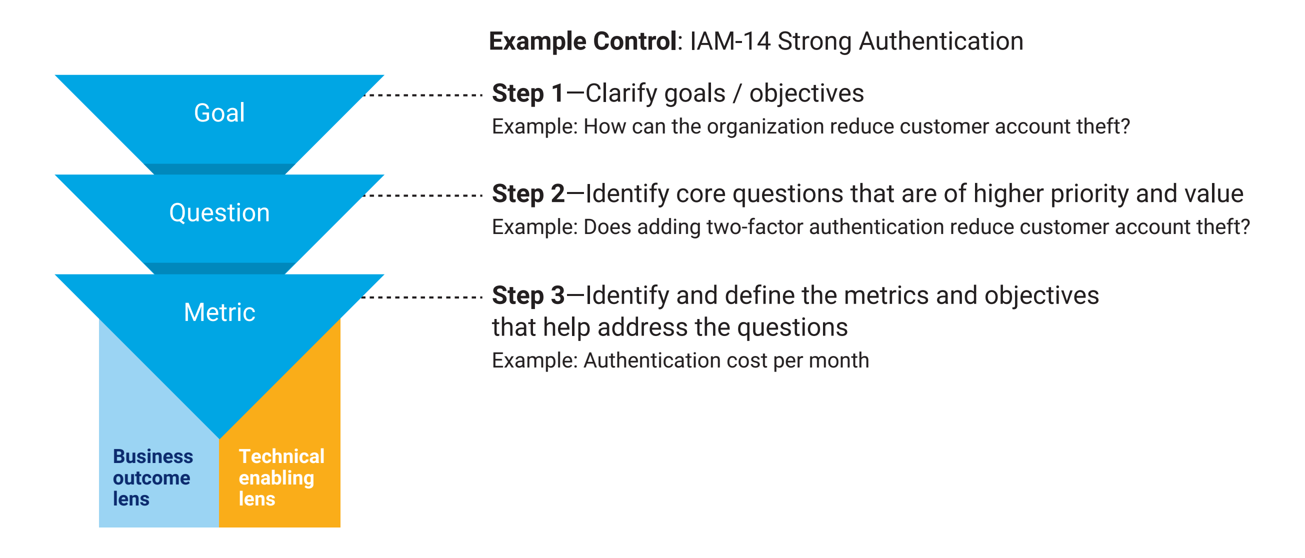

Example Control: IAM-14 Strong Authentication

Step 1—Clarify goals / objectives
Example: How can the organization reduce customer account theft?

Step 2—Identify core questions that are of higher priority and value
Example: Does adding two-factor authentication reduce customer account theft?

Step 3—Identify and define the metrics and objectives that help address the questions
Example: Authentication cost per month

Two major metrics lenses exist: from a business outcome perspective, and from a technical enabling perspective.

For example, improving customer satisfaction is a business perspective. A technical perspective is systems availability. Customer-focused outcome metrics are important to keep a focus on customers and outcomes. The technical-enabling metrics provide other insights to help deliver on business expectations. By measuring, managing and optimizing availability, the business can meet objectives such as customer satisfaction. Many organizations now focus on business outcome metrics; however, it is important to retain technical enabling metrics that provide the early insights for proactive action. The key message is that both business and technical perspectives are required. Having more metrics is not necessarily better—it is important to use the right metrics at the right time to take appropriate action. There are many metrics standards and tools that have a specific focus and value proposition. Using tools and standards that do not align with prioritized goals and objectives lead to failure.

Measurement results always require meaningful context that is provided by a well-defined metric. For example, kilograms are a good unit of measurement. Two-hundred kilograms needs context and answers to other potential questions: What are you weighing and why? Is it the weight of an elephant? How will the information be used—is it just meant to be educational or informational? Is it part of a calculation for lift capacity of a helicopter? Does it influence the amount of feed required? What other measures are required to provide further context and perspectives?

The automated collection of measurement results based on reliable metrics enables the dynamic, proactive allocation of services, while providing the transparency and accountability to build trust.

2.9.4 Cloud Computing Metrics Context

Measurement results are often collected, policies and rules set, and actions defined through setting parameters within cloud products. Measurement and metrics are one of the essential characteristics of the cloud, enabling service usage consumption models, while delivering the resiliency required to build the trust needed for cloud consumption.

In cloud computing, the goals and outcomes are largely influenced by the hosting, deployment and service models. For example, an externally hosted public-cloud SaaS solution has unique goals and metrics that differ from an internally hosted private-cloud IaaS solution. Likewise, defining parameters varies based on these options and on the specific cloud products for which the parameters are defined.

In 2017, the Cloud Security Alliance produced the report, "Improving Metrics in Cyber Resiliency,"[172] which focuses on the metrics of elapsed time to identify failure (ETIF) and elapsed time to identify threat (ETIT). The report includes cyberresilience models that evaluate the relationships of loss and recovery of functionality over time from cybersecurity incidents. The report provides a good example of focusing on just a couple of priority metrics and then performing an interesting analysis of the results provided through those metrics.

Service level agreements (SLAs) are a key subject area reliant on measurement results and metrics. Some contracts and SLAs can contain security attributes that will be measured—e.g., the number of failed logins; or number of users who gained privileged access rights without invoking multifactor authentication. ISO/IEC SC38 has several standards related to cloud computing SLAs and metrics:

- ISO/IEC 19086-1: 2016 *Information technology—Cloud computing—Service level agreement (SLA) framework—Part 1: Overview and concepts*[173]
- ISO/IEC 19086-2: 2018 *Information technology—Cloud computing—Service level agreement (SLA) framework—Part 2: Metric model*[174]

Auditors need to understand the key priorities and goals with the measurement results and metrics that are most applicable to gain the perspectives needed to perform their audit activities. Auditors need to understand the source and context of the measurement results and metrics provided, so that the measurement results and metrics are used (by humans or through automation) to provide assurance and take proactive action. Tools that provide measurement results and metrics can come directly from CSPs, third-party products, existing logs or other information sources.

The available metrics vary based on the competencies, processes, tools and parameters set within the tools. IaaS metrics are typically more mature and widely available. These measures are operational in nature for compute, storage and network. PaaS is typically a little less mature, dealing with development environments. Sadly, SaaS applications are frequently immature, with metrics for specific application services lacking the insights to optimize

[172] Cloud Security Alliance, "Improving Metrics in Cyber Resiliency," 30 August 2017, https://cloudsecurityalliance.org/artifacts/improving-metrics-in-cyber-resiliency/

[173] ISO, ISO/IEC 19086-1:2016 *Information technology – Cloud computing – Service level agreement (SLA) framework – Part 1: Overview and concepts*, September 2016, www.iso.org/standard/67545.html

[174] ISO, ISO/IEC 19086-1:2016 *Information technology – Cloud computing – Service level agreement (SLA) framework – Part 2: Metric model*, December 2018, www.iso.org/standard/67546.html

and effectively manage cloud services. In many cases, auditors need to collect and calibrate the data to gain insights and make informed recommendations.

Software measurement and metrics are significant subject areas, often with complex models and algorithms. Similarly, the analysis and definition of security parameters within complex cloud services can make it difficult to gain transparency and to govern accordingly, using available automation.

See chapter 4 for further discussion and examples related to metrics.

2.10 Cloud Security Certification, Attestation and Validation

The following sections describe the leading types of security certification, attestation and compliance validation reports. In many cases, a CSP or its customer refer to all of these as certifications.

Audit professionals, known for being precise, are the first to point out that inaccuracy, because each carries a specific meaning within the respective standard. However, all of these provide a cloud customer with some level of assurance, comfort or trust that the CSP complies with the applicable standards.

2.10.1 What Is a Cloud Security Certification?

A few years ago, when the cloud computing revolution began, most organizations faced a dilemma. On the one hand, the cloud seemed to offer clear benefits regarding cost and security. On the other hand, it created uncertainty regarding compliance and trust. Since then, this dilemma has often been successfully solved through certification or attestation, based on industry-wide standardized compliance frameworks. CSPs have submitted their services to the scrutiny of the community, through self-assessment results published in public registries, such as the CSA STAR Registry; and independent external auditors, giving their customers a certain degree of assurance and trust.

It is not optimal for a CSP to allow customers to audit its IT operations directly. Instead, the CSP may assign an independent audit and assessment firm to certify or attest to its information security management system and controls.

This section focuses on certification, which is defined as the provision by an independent body of written assurance (a certificate) that the product, service or system in question meets specific requirements.[175]

The current landscape for cloud security certification standards is already quite mature. There are several general security standards and cloud-specific security standards, such as:

- ISO27001 and 27017 (and 27018) certification
- CSA STAR Certification

ISO/IEC 27001

ISO/IEC 27001-2013 is the *Information technology – Security techniques – Information security management systems – Requirements*[176] standard. It specifies the requirements for establishing, implementing, maintaining and continually improving an information security management system within the context of the organization. It includes requirements for the assessment and treatment of information security risk tailored to the organization needs. An accredited certifying body issues a certification upon a successful audit against these requirements. This allows more effective implementation of a process to ensure the confidentiality, integrity and availability of the information and information assets of an organization, minimizing the risk and likelihood of an organization being compromised.

[175] ISO, "Certification," www.iso.org/certification.html

[176] ISO, ISO/IEC 27001 *INFORMATION SECURITY MANAGEMENT*, www.iso.org/isoiec-27001-information-security.html

ISO 27001 enables the demonstration of a commitment to comply with global best practices, proving to customers, suppliers and stakeholders that security is paramount.

The main body of the standard outlines the requirements of a system to manage information security. Annex A is a helpful list of reference control objectives and controls that includes a list of controls for improving the security of information and assets. These controls, along with others that the organization may require, are selected by assessing the risk facing the organization and the applicability of the controls to manage, treat and mitigate those risk conditions. The combination of the controls and the management system to maintain the controls makes ISO/IEC 27001 a highly effective information security standard.

The standard follows a common approach for international systems management standards, making it easy to integrate with other systems the organization might already have in place. The seven core elements of the standard follow:

1. Context of the organization
2. Leadership
3. Planning
4. Support
5. Operation
6. Performance Evaluation
7. Improvement

The information security management system formally defined by ISO/IEC 27001 uses a summary of ISO/IEC 27002 in Annex A to suggest potential information security controls. ISO/IEC 27002:2013 provides guidelines for organizational information security standards and information security management practices, including the selection, implementation and management of controls, taking into consideration the information security risk environment of the organization.

It is designed to be used by organizations to do the following:

- Select controls during implementation of an ISMS based on ISO/IEC 27001
- Implement commonly accepted information security controls
- Develop its own information security management guidelines

However, organizations are free to select and implement other controls as they see fit. In practice, most organizations that adopt ISO/IEC 27001 also use ISO/IEC 27002 as a framework or starting point for their controls, making changes as necessary to suit their information risk treatment requirements. ISO 27001 does not mandate the use of Annex A controls. Other control frameworks can be used. Annex A often is supplemented with control objectives from more specific frameworks, such as ISO 27017, ISO 27018, ISO 27701 and CSA STAR.

CSA STAR Certification

The CSA STAR Certification is a rigorous third-party independent assessment of the security of a CSP. The technology-neutral certification leverages the requirements of the ISO/IEC 27001 management system standard together with the CSA Cloud Controls Matrix, a specified set of control objectives designed to measure the capability levels of the cloud service. The STAR Certification is based on achieving ISO/IEC 27001 and the specified set of criteria outlined in the Cloud Controls Matrix (CCM). STAR Certification evaluates the efficiency of an organization ISMS and ensures the scope, processes and objectives are fit for purpose. It helps organizations prioritize areas for improvement and leads them toward business excellence.

2.10.2 What Is a Cloud Security Attestation?

This section focuses on attestation, which is defined as the issuance of a statement that conveys assurance that the specified requirements have been evaluated (and eventually fulfilled). Such an assurance does not, of itself, afford contractual or other legal guarantees.[177]

While organizations use certification to establish and certify an ISMS at the program level, an attestation is typically designed to focus specifically on a set of commitments or criteria. Attestations are almost always performed by certified public accountants (CPAs, or similar professional designations, depending on the country). Although accounting traditionally focused on financial accounting, audit and tax, the audit of information technology has been the largest growing practice area for CPAs during the past two decades. University programs and even the uniform CPA exam now recognize the critical role of IT in modern commerce and the need for CPAs to be able to audit technology controls.

SOC 1 / SOC 2 / SOC 3

The American Institute of CPAs uses the term service organizations to describe entities that perform services as their main business. This is contrasted with the term user organization, which uses the services. For more than 20 years, CPAs performed specialized audits of IT internal controls, targeting service organizations. During this time, a report by a CPA firm has become the standard for reporting on internal controls at a service organization, as required by the US Securities and Exchange Commission (SEC), the financial services industry, and standard contractual terms with countless service organization customers. One of the main reasons for this wide adoption is that the professional standards that underpin these CPA reports provide customers with a basis for relying on the report conclusions. The objective of these system and organization control reports (SOC) is to provide customers and the auditors of those customers with assurance over the effective operation of IT controls that are designed to address IT risk to information processing. To provide the framework for CPAs to examine controls and to help management understand the related risk, the American Institute of Certified Public Accountants (AICPA) established three SOC reporting options (SOC 1, SOC 2 and SOC 3 reports).

While the SOC reporting based on AICPA Statement on Standard for Attestation Examinations (SSAE) 18 standard is very US-centric, the International Auditing and Assurance Standards Board (IAASB) leverages its International Standard for Assurance Engagements (ISAE) 3000 standard, which is an assurance standard for compliance, sustainability and outsourcing audits designed to address compliance needs beyond the US borders. ISAE 3000 deals with assurance of nonfinancial information. SOC 2 reports, in accordance with certain criteria (trust service principles/sustainability guidelines) that do not impact financial information, should be audited in accordance with the ISAE 3000 standard. Outsourced services that impact financial information of the user organization should be audited in accordance with ISAE 3402. The ISAE 3402 standard is subject to the requirements of ISAE 3000.

To summarize, ISAE and SSAE standards guide auditors. The cloud provider does not comply with any of these. CPAs (or their equivalents) perform SOC reports by following these standards, based on the criteria and objectives provided in the following paragraphs.

SOC 1

All SOC engagements are performed in accordance with SSAE 18. A SOC 1 report is based on AT-C Section 320 of SSAE 18 and is titled “Reporting on Examination of Controls at a Service Organization Relevant to User Entities’ Internal Control Over Financial Reporting”. SOC 1 reports focus solely on controls at a service organization that are likely to be relevant to an audit of the financial statements of a user entity.

[177] ISO 17000:2004, 5.2

Use of a SOC 1 report is restricted to existing (not potential) customers. There are two types of SOC 1 reports:

- **Type 1**—A report on management's description of the service organization system and the suitability of the design of the controls to achieve the related control objectives included in the description as of a specified date.
- **Type 2**—A report on management's description of the service organization system and the suitability of the design and operating effectiveness of the controls to achieve the related control objectives included in the description throughout a specified period.

One of the differentiating features of a SOC 1 report is that it is not based on any defined criteria. The service organization, working with the user entities (i.e., customers), has to define the controls for which it is responsible that would ultimately impact the customers' controls over financial reporting. Most CSP SOC 1 reports list IT general controls, such as logical access, computer operations, change control, and physical and environmental controls as applicable. Additionally, reports may include application controls for processing of financial information, computations of financial metrics and reporting.

SOC 1 SM Reports

Relevant Professional Standards: AT-C Section 320

SOC 2

Recognizing that customers' need for assurance extended beyond financial objectives, the AICPA, in collaboration with Canadian Institute of Chartered Accountants (CPA Canada), formulated the initial version of the Trust Services Criteria (TSC) in 2002 to assist in brokering a trust relationship between the vastly increasing IT service and data processing industry and its customers. The objective was to provide a framework for a CPA to report on the design and operating effectiveness of security, confidentiality, availability, privacy and processing integrity controls.

Unlike SOC 1, SOC 2 engagements use the predefined criteria in the TSC to form the subject matter of the reports. The auditors use the SSAE 18 and ISAE 3000 standards.

Like a SOC 1 report, either a type 1 or type 2 report may be issued. The report provides a description of the service organization system. A type 2 report includes a description of the tests performed by the service auditor and the results of those tests.

SOC 2 reports specifically address one or more of the following five key system attributes, which are aligned with the commitments made to customers:

- **Security**—Information and systems are protected against unauthorized access, unauthorized disclosure of information, and damage to systems that can compromise the availability, integrity, confidentiality and privacy of information or systems, and consequently affect the ability of the entity to meet its objectives.
- **Availability**—Information and systems are available for operation and use to meet the objectives of the entity.
- **Processing integrity**—System processing is complete, valid, accurate, timely and authorized to meet the objectives of the entity.
- **Confidentiality**—Information designated as confidential is protected to meet the objectives of the entity.
- **Privacy**—Personal information is collected, used, retained, disclosed and disposed of appropriately to meet the objectives of the entity.

The service organization selects which of the above attributes are included. The criteria from the SOC 2 TSCs are mid- to high-level and objective-based. There is no requirement for a specified number of characters in a password, or that an intrusion prevention system (IPS) be implemented. The objective-oriented criteria are matched with specific service organization commitments that are made to customers. In that respect, the auditor is not just attesting against a specific standard, but against adherence to criteria based on commitments to customers.

The SOC 2 report was designed with the flexibility to address criteria, such as HIPAA,[178] e-prescribing, NIST CSF, and other IT (security) requirements. The STAR Attestation report is a SOC 2 report complemented with controls from the CCM.

With the rise of cloud computing, there has been a resurgence in demand for reporting by CPA firms on controls related to security, confidentiality and availability. Large CSPs provided or are in the process of providing their customers with SOC 2 reports to address the demand.

SOC 2 SM Reports

Relevant professional standards: AT-C Section 105 and SSAE 18, and TSP section 100, Trust Services Criteria for Security, Availability, Processing Integrity, Confidentiality, and Privacy (2017), otherwise known as the AICPA Trust Services Criteria.

The intended users of the report are management of the service organization and other specified parties who have sufficient knowledge and understanding of the following:

- The nature of the services the organization provides
- How the service organization's system interacts with user entities, subservice organizations and other parties
- Internal controls and limitations
- Complementary user-entity controls and how they interact with related controls at the service organization to meet the applicable trust services criteria
- The applicable trust services criteria
- The risk that may threaten the achievement of the applicable trust services criteria and how controls address that risk.

Unlike SOC 1 reports, prospective customers can use a SOC 2 report.

SOC 3

SOC 3 engagements also use the predefined criteria in the TSC that are used in SOC 2 engagements. The key difference between a SOC 2 report and a SOC 3 report is that a SOC 2 report, which is generally a restricted-use report, contains management's assertion—an often less detailed description of the service—and the service auditor's opinion on the description of the service organization system. A SOC 3 report does not include the detailed controls and results of tests of those controls.

A SOC 3 report is a general-use report that provides only the auditor's report on whether the system achieved the trust services criteria. It contains no description of tests and results or opinion on the description of the system.

SOC 3 SM Reports

Relevant professional standards: AT-C Section 105 and SSAE 18, and TSP section 100, Trust Services Criteria for Security, Availability, Processing Integrity, Confidentiality, and Privacy (2017) otherwise known as the AICPA Trust Services Criteria. Intended users of report: anyone.

[178] Health Insurance Portability and Accountability Act

SOC for Cybersecurity Report

The SOC for Cybersecurity report was developed by the AICPA to allow organizations a means to document and demonstrate the components of their cybersecurity risk management program.

There are a few distinguishing factors for this type of report:

- SOC for Cybersecurity can leverage frameworks other than the TSCs, including but not limited to ISO 27001 or NIST CSF.
- Unlike a Type 1 or Type 2 SOC 2 report, there are no control listings. The report contains the opinion letter, management's assertion, and a description of the cybersecurity risk management program, so in that respect it may look more like a SOC 3.
- SOC for cybersecurity reports can include controls in operation testing.

CSA STAR Attestation

The CSA STAR Attestation is a rigorous third-party independent assessment of a CSP's security. It is based on type 1 or type 2 SOC attestations supplemented by the criteria in the Cloud Controls Matrix (CCM). This assessment:

- Is based on a mature attestation standard
- Allows for immediate adoption of the CCM as additional criteria to AICPA trust service criteria (TSC) and the flexibility to update the criteria as technology and market requirements change
- Does not require the use of any criteria that were not designed for or readily accepted by cloud providers
- Provides for robust reporting on the CSP description of its system and controls, including a description of the service auditor's tests of controls, thereby facilitating market acceptance
- Provides evaluation over time and at a point in time

BSI Criteria Catalogue C5

The Cloud Computing Compliance Criteria Catalogue (C5)[179] defines a baseline security level for cloud computing. It is used by professional CSPs, auditors and cloud customers.

The Federal Office for Information Security in Germany (BSI Germany) initially introduced C5 in 2016. The Criteria Catalogue supports customers in selecting, controlling and monitoring their CSPs. Nationally and internationally established standards form the foundation for the design of the C5 criteria and the requirements for proving conformity, specifically, ISAE 3000—the German Audit Standard (PS) 860 "IT-Prüfung außerhalb der Abschlussprüfung" of the Institut der Wirtschaftsprüfer (IDW), which is in line with ISAE 3000—or other national equivalents to ISAE 3000. Auditors should consider one of these standards or national equivalents as a basis for audit planning, execution and reporting. Auditors should consider further audit standards for individual questions of audit execution and reporting.

These include ISAE 3402 "Assurance Reports on Controls at a Service Organization," the German IDW PS 951 n.F. "Die Prüfung des internen Kontrollsystems bei Dienstleistungsunternehmen," which is in line with ISAE 3402, or other national equivalents to ISAE 3402. Requirements for the contents of the description of the service organization's system, which is part of the audit report, were derived from these standards (cf. Section 4.4.4.1).

[179] Federal Office for Information Security, "New Release C5:2020," www.bsi.bund.de/EN/Topics/CloudComputing/Compliance_Criteria_Catalogue/C5_NewRelease/C5_NewRelease_node.html;jsessionid=CAECB9DB59BF7BE1366E1CF53DDA6F31.2_cid501

In addition, the audit standard AT-C section 105 "Concepts Common to All Attestation Engagements" and AT-C section 205 "Examination Engagements" of AICPA have been taken into account. These standards supplement ISAE 3402 and IDW PS 951, especially with requirements for the consideration of subservice organizations.

The corresponding audit reports build the foundation for a solid risk assessment. In 2019, the Criteria Catalogue was reworked thoroughly, adapting to new developments and increasing quality.

While C5 is based on SOC 2, it adds additional controls that provide information pertaining to data location, service provisioning, place of jurisdiction, existing certification, information disclosure obligations and a full-service description. Using this information, cloud customers can evaluate how regulations (i.e., data privacy), their own policies or the threat environment relate to their use of cloud services.

2.10.3 What Is a Cloud Security Authorization or Validation?

Positioned between certification and attestation, both of which have specific industry terms and requirements, are more purposeful security and compliance authorization and validation assessments. These include FedRAMP, Singapore Multi Tier Cloud Security (MTCS) and PCI DSS, among others. In these cases, a trained and authorized assessor performs a security assessment and issues a report, which authorizes a cloud customer's use of a CSP in a particular context (e.g., hosting government data or processing credit card transactions).

FedRAMP

FedRAMP is a US federal government-wide program that provides a standardized approach to security assessment, authorization and continuous monitoring for cloud-based services. FedRAMP uses a do-once, use-many-times framework that intends to save costs, time and staff, avoiding superfluous agency security assessments and process monitoring reports.

The purposes of FedRAMP follow:

- Ensure that cloud-based services have adequate information security.
- Eliminate duplication of effort and reduce risk management costs.
- Enable rapid and cost-effective procurement of information systems/services for federal agencies.

The FedRAMP assessment process, called a security authorization, is initiated by US federal agencies and CSPs using the FedRAMP requirements, which are FISMA-compliant and based on the NIST 800-53 and 800-37 standards. The FedRAMP Program Management Office (PMO) was established within the General Services Administration (GSA). It is responsible for the development, management and operational support of the FedRAMP program.

Before FedRAMP, vendors were faced with different security assessment requirements for each agency, which meant they had to prepare multiple authorization packages. FedRAMP implemented standard security baselines and processes to provide both an initial authorization of a cloud service and a mechanism for that security package to be reused across the federal government.[180] CSPs must implement the FedRAMP security requirements in their environment and hire a FedRAMP approved third-party assessment organization (3PAO) to perform an independent assessment to audit the cloud system and provide a security assessment package for review.

There are two paths to achieve FedRAMP authorization. The first and most common is an agency authorization. An initial sponsoring agency reviews the assessment package and issues an authority to operate (ATO). The other path involves the FedRAMP Joint Authorization Board (JAB) reviewing the security assessment package based on a prioritized approach, according to the JAB prioritization criteria, and deciding whether to grant a provisional

[180] FedRAMP, "Cloud Service Providers," https://www.fedramp.gov/cloud-service-providers/

authorization. Regardless of the path, federal agencies can leverage CSP authorization packages for review and shared security responsibility model implementation when granting an agency ATO, saving time and money.

FedRAMP uses a shared security responsibility and risk management model that can be leveraged among agencies based on consistent security baselines. FedRAMP provides processes, artifacts and a repository that enable agencies to leverage authorizations with the following:

- Standardized security requirements and ongoing cybersecurity for selected information system impact levels
- A conformity assessment program that identifies certified independent, third-party assessments of security controls implemented by CSPs
- Standardized contractual language to help agencies integrate FedRAMP requirements and best practices into acquisitions
- Repository of authorization packages for cloud services that can be leveraged government-wide
- Standardized shared security responsibility model ongoing assessment and authorization processes for multitenant cloud services

The FedRAMP security authorization process has four distinct areas:

1. **Security assessment**—Security assessments are based on the actions of three partners: CSP, agency CISO or JAB, and 3PAOs. In most cases, the agency works with the CSP in starting the assessment. The CSP must hire a 3PAO to review documentation and conduct penetration testing. Agency CISAs or the JAB reviews the package and makes the decision to grant the ATO to the CSP, or an agency may request a provisional ATO. In either case, it is up to the CSP to submit a security package that is compliant and easy to review. The review process follows the NIST 800-37 risk management framework as tailored for a shared responsibility environment. The CSP identifies the appropriate baseline, implements appropriate security controls and documents the implementation. The CSP contracts with an accredited 3PAO to independently verify and validate its security implementations and its security assessment package. The CSP submits the package either to the agency CISO or the JAB. In the case of JAB submission, the FedRAMP office conducts the review. After documentation is completed and test results are available, the assessment is measured against the FedRAMP requirements to determine if the risk is acceptable. Agencies can grant their own ATOs. The JAB can only grant a provisional authorization. Agencies can then leverage the JAB provisional authorization as the baseline for granting their own ATO.
2. **Authority to operate (ATO)**—The PMO maintains a repository of FedRAMP provisional JAB and agency authorizations and associated security assessment packages for agencies to review. Agencies can use the provisional authorizations and security assessment packages as a baseline for granting their own ATO. If necessary, agencies can add additional controls to the baseline to meet their particular security profile. FedRAMP formally authorizes cloud services through the issuance of formal JAB or agency authorization letters in lieu of a certificate.
3. **Ongoing assessment and authorization (continuous monitoring)**—After a system receives a FedRAMP authorization, it is probable that the security posture of the system could change over time due to changes in the hardware or software on the cloud service offering, or due to the discovery and provocation of new exploits. Ongoing assessment and authorization provide federal agencies using cloud services a method of detecting changes to the security posture of a system for the purpose of making risk-based decisions. The CSP is responsible for continuously monitoring its cloud service offering to detect changes in the security posture of the system to enable well-informed risk-based decision making. The FedRAMP PMO has issued several documents that guide the CSP in performing continuous monitoring.
4. **3PAO accreditation**—CSPs applying for an ATO must use an accredited 3PAO to perform independent security control assessments as required by NIST 800-37 and 800-53A. 3PAOs are accredited by the American Association of Laboratory Accreditation (A2LA), an independent nonprofit that reports to the FedRAMP PMO. The approval process requires applicants to demonstrate their technical capabilities and their independence as an assessor. The approval process follows the conformity assessment approach outlined in ISO/IEC 17020. FedRAMP maintains a list of approved 3PAOs from which CSPs can choose.

Multi-Tier Cloud Security (MTCS)

In 2013, the Infocomm Development Authority of Singapore (now known as Infocomm Media Development Authority of Singapore) and the Singapore IT Standards Committee developed the Multi-Tier Cloud Security (MTCS SS 584) standard. This standard describes the relevant cloud computing security practices and controls for public cloud users, public CSPs, auditors and certifiers. Because security risk requirements differ from user to user, this multitier model specifies different control measures for different levels of security requirements.

MTCS seeks to address needs such as transparency of cloud users. Transparency is a way to build trust between CSPs and cloud users.

With the MTCS, certified CSPs are obliged to spell out the levels of security they can offer to their users. This is done through third-party certification and a self-disclosure requirement for CSPs covering service-oriented information normally captured in SLAs. The disclosure covers areas including data retention, data sovereignty, data portability, liability, availability, business continuous planning/disaster recovery planning (BCP/DRP), and incident and problem management. MTCS SS 584 has three tiers of security, Tier 1 being the base level and Tier 3 being the most stringent.

- **Tier 1**—Designed for non-business-critical data and systems, with baseline security controls to address security risk and threats in potentially low-impact information systems using cloud services (e.g., a website hosting public information)
- **Tier 2**—Designed to address the need of most organizations running business-critical data and systems through a set of more stringent security controls to address security risk and threats in potentially moderate-impact information systems using cloud services to protect business and personal information (e.g., confidential business data, email, customer relation management systems)
- **Tier 3**—Designed for regulated organizations with specific requirements and more stringent security requirements, sometimes in conjunction with industry-specific regulations to supplement and address security risks and threats in high-impact information systems using cloud services (e.g., highly confidential business data, financial records, medical records)

MTCS includes 535 controls. (This may change with version 2 being prepared as of June 2020.)

PCI DSS

The Payment Card Industry Data Security Standard (PCI DSS) is required for any entity that stores, processes or transmits credit card data. It is overseen by the major global card brands and implemented by the PCI Security Standards Council (PCI SSC).

Entities responsible for compliance with PCI DSS are categorized as merchants and service providers. Merchants accept credit card payments for goods and services provided to customers. Service providers are any organization that helps facilitate those transactions, from payment processors and gateways to cloud computing providers. Most CSPs are defined as service providers, meaning they provide some sort of hosting or other service that may be a conduit for credit card data processing. A CSP does not have to store credit card data to be in scope for PCI; it can simply transmit. Additionally, CSPs that impact the security of their customers' cardholder data environments, such as managed security service providers, also fall into scope.

An accredited qualified security assessor (QSA) performs a PCI DSS assessment. There are more than 300 specific control requirements for PCI, including specifications for password controls and multifactor authentication, among others. PCI is generally considered one of the strictest and most prescriptive standards.

The deliverables include a detailed report on compliance (ROC) and attestation of compliance (AOC), which are then shared with customers, banks, the card brands, or other related parties. A cloud provider is deemed validated if it meets 100 percent of the requirements.

2.10.4 The Value and Purpose of a Self-Attestation

CSA STAR Self-Assessment is a due diligence process based on the CSA best practice Consensus Assessments Initiative Questionnaire (CAIQ) and Cloud Controls Matrix (CCM). The CSP voluntarily publishes the results of the self-assessment on the CSA STAR website,[181] which is freely available to the public and open to all cloud providers.

Cloud providers can submit two different types of reports to indicate their compliance with CSA best practices:

- The Consensus Assessments Initiative Questionnaire (CAIQ), which provides industry-accepted ways to document what security controls exist in IaaS, PaaS and SaaS offerings. The questionnaire (CAIQ) provides a set of 310 questions a cloud consumer and cloud auditor may wish to ask of a CSP. Providers may opt to submit a completed Consensus Assessments Initiative Questionnaire.
- The Cloud Controls Matrix (CCM), which provides a controls framework that provides a detailed understanding of security concepts and principles aligned to the Cloud Security Alliance Guidance in 13 domains. As a framework, the CSA CCM provides organizations with the necessary structure, detail and clarity relating to information security requirements tailored to the cloud industry. Providers may choose to submit a report documenting compliance with the Cloud Controls Matrix.

2.10.5 The Compliance Fatigue Problem

Security compliance based on third-party audit is becoming increasingly complex, especially due to the considerable number of national, international and industry-specific standards and certification schemes present in the market.

Proliferation of Certification Schemes

The proliferation of certification schemes is generating compliance fatigue in the industry—not to mention contradicting audit reports related to similar controls, especially for companies operating on a global scale. That often translates into substantial costs for the service providers that can afford compliance with multiple standards; potential barriers to market entry for smaller providers; and confusion for users/customers, who are not necessarily experts on certification and standards, and who might have trouble in understanding which compliance seal to rely on.

The Value of a Multi-Standard Compliance and Auditing Program

At first sight, all these new certification schemes seem to be uniquely heterogeneous, because they target wider or specific application areas (e.g., national, sectoral, and regulatory domains and requirements), but this is not the case. Cloud-based certification schemes are based on globally accepted and widely used standards (e.g., ISO/IEC 27001). Consequently, their very core security domains and requirements are largely homogeneous from the perspective of security requirements and objectives. As a result, many of the existing cloud security standards, especially national ones, include many overlapping requirements. It is valuable to identify common denominators between these requirements and present them under a comprehensive framework.

EU-SEC,[182] a European Union-funded project, conducted a study of compliance schemes. Researchers observed that many certification schemes' individual security requirements and control objectives are largely the same. Most of the

[181] Cloud Security Alliance, "CSA Security Trust Assurance and Risk (STAR)," https://cloudsecurityalliance.org/star/
[182] EU-SEC, "The European Security Certification Framework (EU-SEC)," https://www.sec-cert.eu/

requirements (71 percent) were satisfactorily covered with the CCM controls, while only 29 percent of the requirements were covered partially (16 percent), or were not covered (13 percent) by CCM. Consequently, when a CSP obtains a certification or attestation under two different schemes, a lot of work is duplicated, unduly increasing costs and complexity. Therefore, it seems that in many cases, the work done under one compliance scheme should be reusable under another. This would allow CSPs to focus on the differences in security requirements between multiple compliance schemes. The mutual recognition approach would streamline the cloud compliance process, bring efficiency, increase assurance and reduce reassessment costs.

It should not be a surprise that most CSPs undergo more than one type of assessment based on the markets they are in. Another trend in the marketplace is the use of custom control frameworks: CSPs identify a standard set of controls that are common across the different compliance frameworks. To be clear, this requires work, and it is not possible to purchase or download a set of controls that guarantees compliance with everything. The CCM is a great starting point for such an initiative, but it needs to be customized and reflect the cloud providers' unique scope of services, responsibilities, and commitments to their customers.

2.11 Chapter 2 Knowledge Check

REVIEW QUESTIONS

1. What is the objective of a cloud compliance program? (Select the correct answer.)

   A. To ensure all aspects related to the cloud adoption in an organization adhere to the requirements for which it is accountable.
   B. To determine the risk appetite and risk tolerance of the organization.
   C. To drive decisions for the measures put in place by the organization.
   D. To validate the organization is aligned with their competition and can remain competitive.

2. What is the difference between legacy and cloud compliance programs?

   A. There is no difference.
   B. Cloud compliance relies on systems with defined boundaries.
   C. Legacy compliance programs have complete control and responsibility for the infrastructure and all requirements that keep it secure.
   D. Cloud compliance programs rely on third parties for the delivery of technology with complete control and influence over the quality, availability and reliability of the service being provided.

3. What is a cloud governance and strategy audit?

   A. An audit that evaluates the framework of the organization for defining requirements, performing risk assessments, monitoring controls, reporting adherence and developing a strategic plan for the cloud.
   B. An audit that covers logging, monitoring, scanning and alerting of a system, account or environment, and can be achieved using real-time automated scripts or manual testing.
   C. A review of all user and service accounts, and permission within your information system boundaries, including on-premises systems, cloud environments and other applications.
   D. An audit performed against technical, administrative and/or physical controls as defined in the organization policies and procedures.

4. Select **ALL** the correct statements concerning legal and regulatory requirements, standards and security frameworks in the cloud:

   A. They serve to guide the cloud controls program of an organization.
   B. They should be monitored to ensure that changes in the requirements are promptly reflected in the cloud compliance program.
   C. They are important for when designing the cloud compliance program, but not so relevant during the execution phase.
   D. The majority of the laws and regulations are cloud specific.

5. In a DevOps IT environment, the following team is responsible for mitigating and/or managing risk on a daily basis:

   A. Internal Audit
   B. Risk Management
   C. DevOps
   D. Quality Assurance

6. Due to the shared responsibility model, control operationalization in the cloud (select all that apply):

   A. Is different from what it is in traditional on-premises IT environments.
   B. Involves a combined effort from multiple parties.
   C. Is simpler than control operationalization in traditional on-premises IT environments.
   D. Should replicate the control environment of traditional on-premises IT environments.

7. What is a common way to come up with metrics that support a decision process?

   A. Apply the Goal, Question, Metric approach.
   B. Apply the Metric, Measurement, Decision feedback loop.
   C. Make sure that the SLA contains automated metrics.
   D. Use the ISO/IEC 19086 standard.

8. What is FedRAMP?

   A. A US federal government program that provides a standardized approach to security assessment of cloud-based services.
   B. A US federal government program that provides a risk management model that can be leveraged among US agencies for IT certification.
   C. A US federal funding program that helps innovative startups in information security achieve ramping growth.
   D. A collaboration between AICPA and the Canadian Institute of Chartered Accountants (CPA Canada) for a federated approach to attestions.

Answers on page 176

Chapter 2 ANSWER KEY

Correct answers appear in **bold** font.

1. **A. Knowing what the organization is being held accountable for is critical to build the compliance program around so each objective is being met.**

   B. Establishing the risk appetite and risk tolerance of the organization are not the objective of a cloud compliance program but, important aspects when designing your cloud compliance program.

   C. The purpose of Enterprise Risk Management (ERM), governance policies, the InfoSec Department, data privacy, etc., is to ensure that decisions are made correctly and have a proper foundation when choosing what to do.

   D. This is not the objective of a cloud compliance program.

2. A. The primary difference between legacy and cloud compliance programs are the control of the service being provided and responsibilities that belong to the organization.

   B. Legacy compliance programs rely on systems with defined boundaries.

   C. Cloud based environments will follow the shared responsibility model where the cloud vendor and customer will document each party's responsibilities. In legacy environments, there is no cloud vendor and the organization will be required to maintain all aspects of the infrastructure.

   D. Cloud compliance programs do rely on third parties for the delivery of technology but have limited control and influence over the quality, availability and reliability of the service being provided.

3. **A. This is the definition of cloud governance and strategy audit provided in the Certificate of Cloud Auditing Knowledge Study Guide.**

   B. That's the definition of Configuration and Activity Monitoring.

   C. That's the definition of 'access review'.

   D. That's the definition of Compliance and Controls Audits.

4. **A. Legal and regulatory requirements, standards and security frameworks are indeed between the main driver of an organization compliance program.**

   B. Legal and regulatory requirements, standards and security frameworks shall be monitored to ensure alignment over time.

   C. Legal and regulatory requirements, standards and security frameworks are relevant during the whole cloud compliance program lifecycle (from design to implementation).

   D. The majority of laws and regulations are not cloud specific, but cloud relevant.

5. A. Internal Audit is a third line of defense team that periodically validates that risk management activities are appropriate, and that controls are in place work as expected

   B. Given the frequency of changes and deployments to production in a DevOps environment, the risk management function doesn't manage risk on a daily basis but checks afterwards if policies and procedures were followed

   C. Given the frequency of changes and deployments to production in a DevOps environment, the DevOps team is responsible to mitigate and/or manage risk on a daily basis

   D. A quality assurance team is not responsible for managing risk

6. **A. Organizations have full control of processes implemented in their traditional on-prem IT environments, while for cloud environments, the Cloud Service Provider (CSP) manages a portion of the technology stack.**

   B. Both the Cloud Service Provider and cloud customer are responsible for implementing and operationalizing controls in the cloud.

   C. Some controls in the cloud may be simpler to implement given the tools the Cloud Service Provider (CSP) makes available to its customers, but other controls may not be possible to implement given that the technology infrastructure in a cloud environment is shared.

   D. As a general rule, It’s important to avoid the temptation of replicating traditional IT practices in the cloud when better cloud-native approaches exist.

7. **A. Correct. This is the proposed approach described in the course.**

   B. Incorrect. This was not suggested as a correct approach.

   C. Incorrect. Automated metrics and measurements are great, but that does not mean they support a decision process.

   D. Incorrect. This standard deals with metrics, but not decision making.

8. **A. As the FedRAMP website states “FedRAMP simplifies security for the digital age by providing a standardized approach to security for the cloud.”**

   B. That’s not the main purpose of FedRAMP.

   C. No, FedRAMP is not a funding program.

   D. No, this is unrated to SOC.

Page intentionally left blank

Chapter 3:

Introducing CCM and CAIQ

Introducing CCM and CAIQ

Introducing CCM and CAIQ

3.1 Learning Objectives

After completing this chapter, learners will be able to:

1. Identify the CSA Cloud Controls Matrix (CCM) and CCM domains.
2. Explain the Consensus Assessment Initiative Questionnaire (CAIQ).
3. Outline CCM and CAIQ structures.
4. Recall CCM relationship with other frameworks (mapping and gap analysis).
5. Compare transition changes from CCM V3.0.1 to CCM V4.

3.2 Overview

Chapter 3 focuses on the CSA Cloud Controls Matrix (CCM) and the Consensus Assessment Initiative Questionnaire (CAIQ). This chapter explains their purpose, why they were created, their objectives and the relationship between CCM and CAIQ. It also covers their connection with the CSA Security Guidance (hereafter referred to as the Guidance) and the CSA Enterprise Architecture Model, together with their role within the context of the Security, Trust, Assurance and Risk (STAR) Program. It gives an overview of their target audience and how each audience is meant to use CCM and CAIQ. It provides guidance on the mappings of CCM controls to other industry-accepted security standards, regulations and frameworks, with examples. This chapter introduces the 16 CCM domains, the 133 control objectives of CCM v3.0.1 and their relationships with the questions of the CAIQ.

3.3 The CCM and How It Was Created

The Cloud Controls Matrix (CCM) is a security control framework that is dedicated to managing risk in the cloud. It was developed to address the lack of a cloud-relevant security and compliance framework, because, at the time of its creation, risk management was solely focused on addressing traditional computing infrastructure. The CCM is meant to fill this gap, by helping guide both cloud service providers (CSPs) and cloud customers assess the overall security risk of a cloud service.

For CSPs, the CCM is meant to establish best practices to support the secure implementation of cloud infrastructure and services. For cloud customers, the purpose is to enable them to better evaluate and assess CSPs.

The CCM includes detailed security concepts and principles that are aligned with the CSA Security Guidance v4.[183] This document provides guidance on how security principles and criteria should be implemented in a cloud architecture, whereas the CCM provides guidance on *what* should be done. The CSA encourages organizations to use the CCM as a companion to the CSA Security Guidance, because it allows the user to identify security controls and understand how they should be implemented.

3.3.1 How It Was Created

The CCM is both vendor-independent and consensus-driven. Like many other frameworks, CCM is the result of the contribution of security experts who volunteer for the CCM Working Group (WG). Since its initial publication in 2010, more than 300 experts have contributed to versions of CCM, and more than 500 have participated in the peer review process. Within the CCM Working Group, the decision-making process is mainly based on consensus, which

[183] Cloud Security Alliance, "Security Guidance v4.0," 26 July 2017, https://cloudsecurityalliance.org/research/guidance/

means that each control objective reflects the shared and common opinion of the experts on a specific technical aspect of cloud security.

After the building process is finalized, the CCM goes through an open peer review process, which is part of the standard CSA Research Lifecycle.[184] The CCM is very resilient, because it is an open framework subject to public scrutiny. Anyone who has the expertise can review and critique it during open peer reviews, making it a very resilient framework. At the same time, the process of collecting input remains open at any time. Because the CCM is a free framework, anyone can download and comment if they see areas for improvement.

The CCM change process is triggered by technology evolution, process improvement, or the introduction of new laws, regulations or technical frameworks. The change process can be triggered by the CSA internal stakeholders, working group members, CSA Corporate Members, CSA chapters or external experts. Change requests are evaluated by the working group and might be accepted.

CCM Release Overview

There are three forms of CSA CCM releases.

A full or major release (i.e., V3.0 → V4.0) entails the following:

- A revision of the CCM domain structure
- A technical change (revision, addition or deletion) of a number of controls greater than 10 percent, compared to the previous full release

A dot release (i.e., V1.0 → 1.1) constitutes the following forms of CSA CCM modifications:

- A change in the matrix column
- A technical change of number of controls smaller than 10 percent

A minor release (i.e. V3.0 → V.3.0.1) constitutes the following forms of CSA CCM modifications:

- **An editorial change in the controls**
- **Updated control framework mapping**—Revisions to previously published mapped control frameworks
- **New control framework mapping**—Introduction of new control frameworks mapped to previously published CSA CCM controls.

3.3.2 CCM Versions

Version one of the CCM was created in 2010. There were several minor updates made to this version until 2014, when there was a major update from version one to version three. Version two was skipped because CSA had released version three of the Security Guidance in 2013 and wanted to align the CCM with the new guidance.

At the end of 2014, CSA released version 3.0.1. This is the standing version[185] with some minor revisions and additional mappings.

A comparison of version one and version three shows the following new domains were added to improve the alignment between the CCM and the Guidance:

- Mobile security
- Supply chain management, transparency, and accountability
- Interoperability and portability

[184] Cloud Security Alliance, "Research Lifecycle," https://cloudsecurityalliance.org/research/lifecycle/
[185] As of October 2020

3.3.3 CCM Target Audience, Purpose and Objectives

Regardless of the type of organization (CSP vs. customer), the nature of its business, its size (large corporation vs. small company), or cloud delivery model (IaaS vs. PaaS vs. SaaS), the CCM can be used to define, implement and enforce its security requirements and monitor their implementation. The CCM assists companies in translating their internal organizational, operational and legal stipulations into a standardized set of cloud-relevant policies, procedures and technical control objectives.

The CCM is also a tool for internal and external assessments or audits. It is designed to be used in alignment with the Consensus Assessments Initiative Questionnaire (CAIQ), which provides a set of yes or no questions that can be answered to determine if the CCM controls are being met. Both documents help auditors understand if an organization is following its internal governance policies and fulfilling its legal and regulatory obligations.

For example, an organization, based on its internal risk assessment, might identify the need to protect the confidentiality, integrity and availability of information related to the manufacturing process. The datasets have different levels of sensitivity and criticality; they are stored in a cloud database and are processed in several cloud-based applications. The organization can use the CCM to identify specific policy, procedural and technical requirements, and define control objectives that will be included in the organizational security program. It uses those control objectives to enforce the requirements vis-à-vis the internal users, the business partners and the CSPs, and to monitor adherence to both internal policies and external compliance requirements.

CCM Target Audience

The CCM was created to help cloud customers, cloud service providers, and auditors and consultants.

Cloud Customers

The CCM allows cloud customers to build a detailed list of requirements and controls they want their CSP to implement as part of their overall third-party risk management and procurement program. It also helps normalize security expectations, provide a cloud taxonomy, and improve understanding of the security measures implemented in the cloud supply chain. Because the actors within a cloud supply chain are independent organizations, each has its own way of expressing and representing its security requirements. It might use a different vocabulary or apply policies that differ from others. In such a context, it is of paramount importance to define a taxonomy, or a set of agreed upon terms, to normalize the different languages. That is why CCM plays a key role, and why more overarching frameworks are necessary to simplify interoperability.

Cloud customers can use the CCM controls to do the following:

- Map organizational, operational and legal requirements to control objectives.
- Build an operational cloud risk management program.
- Build a third-party risk management program.
- Build an internal and external cloud audit plan.

When an organization is building a cloud risk management program, the CCM can help it measure, assess and monitor the risk associated with a CSP or a particular service. It allows a customer to understand the gaps between its own security needs and the CSP security capabilities. The customer can then use it to identify the compensating controls to close the gap between the organization's needs and the provider's offerings.

When building a third-party risk management program, the CCM allows customers to assess a cloud service during the overall service life cycle. It can be used to make the initial evaluation of the service before its acquisition, compare offerings from different CSPs, and monitor alignment with internal requirements during the service execution.

Cloud Service Providers (CSPs)

The CCM serves multiple purposes for CSPs. First and foremost, it offers cloud-specific, industry-validated best practices they can follow to guide their internal security programs. In addition, it offers standardized language that the CSP can use to communicate with its customers and business partners.

The CCM mapping feature provides the opportunity to show alignment with other recognized international, national and industry frameworks, and compliance with the CSA STAR program, which relies on CCM as one of its foundational frameworks (see chapter 9 for more details). The CSA STAR program allows organizations to be more transparent and reduce the number of security questionnaires they must provide for individual customers. One way they can do this is by completing the CCM extended question self-assessment—the Consensus Assessment Initiative Questionnaire (CAIQ)—and submitting it to the CSA STAR Registry, a free, publicly accessible registry that documents the security controls provided by CSPs.

CSPs can use the CCM controls to do the following:

- Build an internal security program based on mature and industry recognized best practices
- Facilitate communication and interoperability with business partners and customers
- Demonstrate commitment to security and transparency about its security posture
- Streamline compliance by leveraging the mapping between CCM controls and the controls in other international, national and industry frameworks
- Reduce the time and effort spent in addressing customer questionnaires
- Demonstrate commitment to security to regulators by adhering to the CSA STAR program (see chapter 9)
- Build a cloud internal and external audit plan

Auditors and Consultants

Auditors and consultants can use the CCM to guide their clients in designing, planning and executing activities dedicated to cloud customers and CSPs.

Consultants and auditors can leverage CSA to do the following:

- Help organizations assess their cloud maturity
- Establish controls aligned with the CCM
- Compare organizations with market peers through benchmarking

3.3.4 Key Features of the CCM

This section describes key features of the CCM.

Mappings to Other Industry-Accepted Standards, Regulations and Frameworks, Internal Control Programs and Control Reports

An important aspect of the CCM is that it maps to other security standards, regulations and frameworks. When the CCM was created, there were already a number of other information security standards, best practices and regulations in existence (e.g., ISO/IEC 27001 and 27002, PCI DSS, NERC CIP, BITS, BSI). Many companies already had their internal structures and frameworks set up and aligned with those standards.

The CSA wanted to provide cloud sector-specific controls and also make sure that organizations had a clear path to connect their existing control frameworks and programs with the new cloud-relevant controls included in the CCM. It, therefore, built all the controls created in the CCM as an extension of existing framework controls. CSA did this by creating a mapping, or a linkage, between a control in a framework, e.g., ISO 27001, and the CCM. The CCM builds on top of the framework to provide a control that is specific to the cloud sector. It then takes this one step further by ensuring that controls are linked to a specific area within a cloud architecture. It helps to identify if a specific control is relevant for IaaS vs. PaaS vs. SaaS, for example. Because the CCM creates links through mapping, it provides an initial internal controls system that identifies which controls should be put in place to further the cloud journey and implementation of an organization.

Another important aspect of the CCM is that it provides direction for organization control reports. For instance, if an organization needs to demonstrate to its board of directors, authorities, regulators, or internal/external auditors how it adheres to specific security requirements, it will need to structure a report according to certain security principles. The value of the CCM is that it already provides a structure for reporting.

See section 3.7 for more details about the CCM mapping to other frameworks, best practices and regulations.

Connection With the CSA Enterprise Architecture Model

The CSA Enterprise Architecture (EA) is a high-level conceptual model that includes a methodology and a set of tools. It enables security architects, enterprise architects and risk management professionals to assess the status of their internal IT and cloud providers in terms of security capabilities, and it helps them create a road map to meet the security needs of their business. The CSA EA identifies a comprehensive set of functional capabilities and processes grouped in domains. The actions included in each domain are based on best-practice architecture frameworks, for example:

- **Sherwood Business Security Architecture (SABSA)**—Defines security capabilities from a business perspective.
- **Information Technology Infrastructure Library (ITIL)**—Defines the capabilities needed to manage the IT services of the company, and thus the security capabilities necessary to manage those services securely.
- **Jericho Forum**—Defines technical security capabilities that arise from the reality of traditional data center technology environments, shifting to one where solutions span the Internet across multiple data centers, some owned by the business and some used purely as an outsourced service.
- **Open Group Architecture Framework (TOGAF)**—Provides an enterprise architecture framework and methodology for planning, designing and governing information architectures, and thus provides a common framework to integrate the work of the security architect with an enterprise architecture.

The Cloud Controls Matrix creates a connection with the CSA Enterprise Architecture or any enterprise architecture model by making sure that the controls are linked to a specific area within a cloud architecture.

Because the CCM is mapped to existing security controls specifications from various legal and regulatory frameworks, and because that same matrix is mapped to the security capabilities of the architecture, it is easy for a company to assess which capabilities are in place for compliance with applicable regulations and best practice frameworks.

Foundation for the CSA STAR Program

Within the CSA ecosystem the CCM provides the foundational framework for the Security, Trust, Assurance and Risk (STAR) program by outlining security controls that a provider should meet to comply with the STAR program. See chapter 9 for more details about the role of the CCM in the STAR program.

Connection With Existing Information Security Control Environments

The CCM provides structure and clarity around cloud security for all organizations regardless of their size or level of maturity, as follows:

- **Organizations new to cloud with mature security programs**—These are organizations that are already adhering to non-cloud specific frameworks and have implemented internal security controls to satisfy their needs and requirements. For them, the CCM provides the required cloud perspective and understanding of how to connect the existing control environment with the newly implemented cloud services.
- **Organizations with a cloud-first strategy**—For these enterprises, the CCM provides guidance on how to increase the level of maturity and integration between their on-premises and cloud control environments.
- **Smaller businesses or startups**—For organizations looking for guidance on how to build a cloud security program, the CCM provides a structure to use without reinventing the wheel. Rather than trying to discover which security requirements pertain to them, they can use the CCM controls.

Control Ownership

The CCM clearly delineates control ownership. One of the key differences between on-premises infrastructure and cloud computing is that the cloud creates the need for a shared responsibility model. The CCM is structured so that it assigns the responsibility for controls implementation to either the CSP or the cloud customer, or both. This is especially important since one of the most common mistakes in cloud computing is not fully understanding the shared responsibility model.

Cloud Service Model Applicability of the Controls

The responsibilities for implementing security controls in cloud computing largely vary depending on the service model (IaaS, PaaS, SaaS). See chapter 1 for more on this topic. It is also covered in the Certificate of Cloud Security Knowledge (CCSK) course.

Some controls clearly are the province of the IaaS providers (for example, data center security-related controls), whereas other controls are applicable to any service model (e.g., identity and access management). The CCM defines the relevance of each control to the three cloud delivery models, helping users to understand what is relevant, and when.

Security Expectations, and Security Measures Implemented in the Cloud Supply Chain

Cloud computing has a complex supply chain with many actors working together. A cloud customer might have multiple IaaS providers and hundreds (or even thousands) of SaaS providers. For all actors in the supply chain to understand one another, everyone needs to have a standard language that uses the same vocabulary. The CCM meets this need by providing a standard taxonomy.

Note: The CCM provides a lingua franca, a language understandable by everyone. This is especially important in a complex supply chain like the cloud, with multiple different actors involved and potentially multiple different sets of security controls.

3.3.5 Relationship Between CCM and the Security Guidance

This section describes the relationship between the CCM and the Security Guidance.

What Is the Security Guidance?

The flagship CSA document is the CSA Security Guidance for Critical Areas of Focus in Cloud Computing v4.0.[186] This document acts as a practical, actionable road map for individuals and organizations that want to safely and securely adopt the cloud paradigm, making it a good starting point for professionals who want to learn more about cloud security considerations.

Use of the CCM vs. the Security Guidance

The Guidance provides a set of best practices and recommendations for how to secure the cloud, whereas the CCM provides a set of control objectives against which an organization should assess cloud security. The Guidance essentially explains how to do things in the cloud, and the CCM explains what to do by identifying specific controls.

The Guidance is structured in domains, similar to the CCM, and focuses on many of the same key areas of cloud computing. Fourteen domains cover the key concepts of cloud computing and architecture. Higher-level domains address cloud governance and enterprise risk management. There are domains for legal and regulatory compliance, key discovery, compliance audit management, and governance. Following the high-level domains, the Guidance drills down to more technical domains, such as management plane and business continuity, infrastructure security and incident response.

The Difference in Domain Structure

There is a substantial alignment between the Security Guidance (14 domains) and the CCM (16 domains). However, the alignment between the two documents is not perfect for several reasons. First, some guidance domains, such as SecaaS, are not suitable controls domains. Also, CSA wanted to be sure that the CCM reflected the body of knowledge gained from working groups created after the Guidance document was released.

Domains included in CCM V3 that do not map directly to the Security Guidance follow:

- **Mobile security**—CSA had set up a new working group to research mobile security after the Guidance was published and wanted to make sure its findings were included in V3.
- **Supply chain management, transparency and accountability**—CSA considered it important to include a specific reference since these are key differentiators between on-premises and cloud computing.

3.4 CCM Domains

This section describes the CCM domains.

[186] *Op cit* Cloud Security Alliance, “Security Guidance v4.0”

3.4.1 Control Domains

The CCM V3.0.1 is structured into 16 security domains and 133 controls (**figure 3.1**). The 16 domains were based on the Guidance document and inspired by major frameworks, such as ISO/IEC 27001 and ISO/IEC 27002. The domain defines what category the controls fall under. CCM was deliberately designed like existing noncloud leading information security frameworks to leverage familiarity with those existing frameworks.

Figure 3.1—CCM V3.0.1 Structure

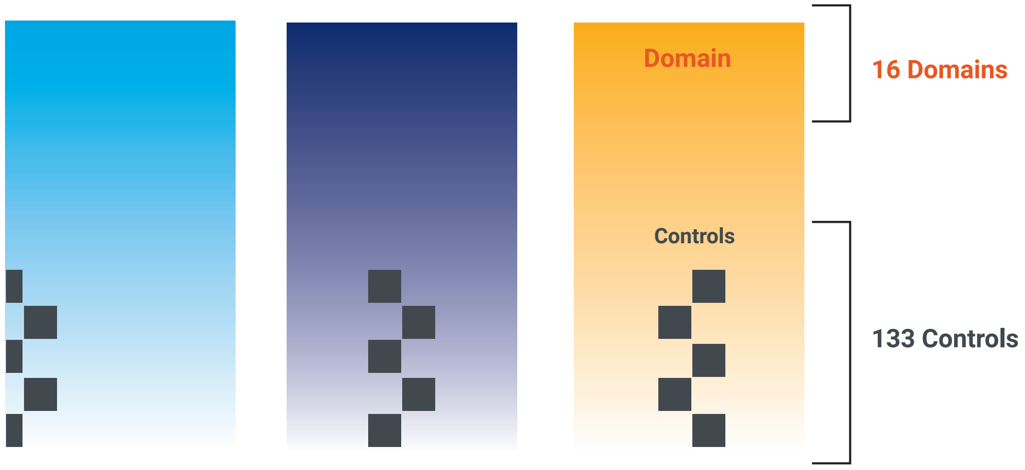

The CCM Security Domains are outlined in **figure 3.2.**

Figure 3.2—CCM V3.0.1 Structure

AIS	Application & Interface Security	HRS	Human Resources Security
AAC	Audit Assurance & Compliance	IAM	Identity & Access Management
BCR	Business Continuity Mgmt & Op Resilience	IVS	Infrastructure & Virtualization
CCC	Change Control & Configuration Management	IPY	Interoperability & Portability
DSI	Data Security & Information Lifecycle Mgmt	MOS	Mobile Security
DCS	Datacenter Security	SEF	Sec. Incident Mgmt, E-Disc & Cloud Forensics
EKM	Encryption & Key Management	STA	Supply Chain Mgmt, Transparency & Accountability
GRM	Governance & Risk Management	TVM	Threat & Vulnerability Management

Source: Cloud Security Alliance, Cloud Controls Matrix (CCM) Working Group

3.4.2 Application and Interface Security

The application and Interface security domain focuses on two key components of the cloud ecosystem: software applications and their user interfaces.

Application security covers the design, threat modeling, development, delivery, implementation and maintenance of production applications. This is a rapidly evolving area (see chapter 8 on DevOps). Cloud computing is one of the biggest drivers for these changes from the technical and organizational perspective, resulting in the need for a more mature approach to application security.

Cloud computing providers expose a set of software interfaces, or application programming interfaces (APIs) that customers use to manage and interact with cloud services. An API is a set of routines, protocols and tools for building software applications. A good API makes it easier to develop software by providing many of the building blocks. A programmer then puts the blocks together.

A large percentage of data exposures are the result of attacks at the application layer, particularly for web applications. Most large data breaches are a result of poor application security. Cloud-based software applications require a design rigor similar to an application connecting to the raw Internet. The threats in a cloud environment will be more numerous than those experienced in a traditional data center.

APIs are a major concern, because security, privacy and availability of general cloud services are dependent upon the security of basic APIs. Provisioning, management, orchestration and monitoring are all carried out using these interfaces.

Organizations and third parties often build upon APIs to offer value-added services to their customers. This increases risk, as organizations may be required to relinquish their credentials to third parties to enable their agency. For example, insecure APIs were blamed for successful attacks on Pinterest and Instagram. API vulnerability also played a role in the breach at messaging firm Snapchat, which exposed the phone numbers and names of up to 4.6 million users.[187]

Some of the major API attack vectors include bypassing authentication defenses, bypassing data validation via third-party APIs, evading detection of brute force authorization, evading rate limits, and abusing content types.

The CCM Application and Interface Security domain is mainly a reflection of the Application Security Domain of the CSA Guidance V4.

Following are controls**:**

- **AIS 01**—Application security
- **AIS 02—**Customer access requirements
- **AIS 03**—Data integrity
- **AIS 04**—Data security/integrity

3.4.3 Audit, Assurance and Compliance

Audit, assurance and compliance have always been critical areas in information security. Their importance is emphasized in chapter 1and chapter 2 of this study guide. Having a clear understanding of compliance requirements and obligations and the level of assurance expected or demanded from customers is of paramount importance for CSPs. Failing to implement an auditing plan can have a substantial negative impact in terms of loss of control and

[187] Collins, K.; "Snapchat hack leaves phone numbers of 4.6 million users partially exposed," Wired, 2 January 2014, https://www.wired.co.uk/article/snapchat-hacked

governance, lack of compliance, and ultimately financial, operational and reputational damages. The focus of compliance is alignment with internal policies and external requirements (e.g., law, regulation, industry frameworks). Examples of regulatory compliance laws and regulations include the EU's GDPR, PCI DSS, HIPAA, and the Sarbanes-Oxley Act (SOX).

Organizations—especially large ones—are expected to comply with an increasing number of industry frameworks. However, frameworks are often based on the same underlying principles (except ones like HIPAA or HITECH, which are specifically focused on healthcare data). If an organization effectively implements the requirements of some of these frameworks, it will fulfill many of the requirements found in other frameworks as well. For example, physical security (such as controlled access points for data centers) is a requisite in most of these frameworks, and if the CSP follows ISO/IEC 27001 or PCI DSS, the requirement will be fulfilled. Given this, the policies should be designed and implemented in such a way that compliance is not an inhibitor of organizational effectiveness, but a complement to internally determined policies.

Cloud customers, especially small organizations leveraging cloud services, have the potential advantage of inheriting compliance from the cloud provider, limited to the scope of service. Even so, the customers remain accountable for their compliance and should always perform their due diligence by reviewing the results of third-party audits and examining the value of existing certifications and attestations.

As the number of rules has increased, regulatory compliance has become more prominent in a variety of organizations. The trend has led to the creation of corporate, chief and regulatory compliance officer roles whose sole focus is to make sure the enterprise conforms to stringent, complex legal mandates.

The CCM Audit, Assurance and Compliance domain is mostly a reflection of the Compliance and Audit Management domain of the CSA Guidance V4.

Following are the controls:

- **AAC 01**—Audit planning
- **AAC 02**—Independent audits
- **AAC 03**—Information system regulatory mapping

3.4.4 Business Continuity Management and Operational Resilience

Traditionally, the three tenets of information security are confidentiality, integrity and availability. Business continuity deals with the availability component of those three requirements. It refers to the capability of the organization to continue delivering products or services at acceptable predefined levels following a disruptive incident (ISO 22301:2019).[188]

Business continuity is often described as common sense. It is about taking responsibility and enabling the organization to stay on course regardless of the storms it is forced to weather (i.e., the ability to keep calm and carry on). This means building and improving resilience in an organization. It starts with identifying key products and services and the most urgent activities that underpin them. Once that analysis is complete, it is about devising plans and strategies that will enable the organization to continue business operations and make a quick and effective recovery from any type of disruption—whatever its size or cause.[189]

Business continuity has often been hailed as one of the biggest benefits of cloud computing, given the very nature of the cloud and the possibility of quickly scaling up and down. However, the concept of IaaS always being available hides substantial risk.

[188] ISO, ISO 22301:2019 *Security and resilience – Business continuity management systems – Requirements*, 2019, www.iso.org/standard/75106.html
[189] BCI, "Introduction to Business Continuity," www.thebci.org/index.php/resources/what-is-business-continuity

The fact that many cloud services have a track record of being offline only for a few hours a year or so is testament to their overall reliability, but customers should not become complacent. Even the biggest and most reliable computer services in the world are not infallible. Over the past few years, there have been several examples of IaaS provider outages that forced their customers offline.[190, 191]

For a cloud customer, the transition to a CSP must include an assessment of the uptime the provider contractually commits to. However, the service level agreement (SLA) may not be enough to satisfy the customer.

Cloud customers need to factor in that a cloud-based infrastructure will experience downtime at some point, and there needs to be a failover mechanism of some kind to cope with it. Customers need to build operational resilience into systems and assume that a failure of some IT component in the application stack is inevitably going to happen at some point.

Although many recommendations focus on documented assertions that a service will maintain continuity, the true test of these assertions takes place during a significant incident. Without waiting for an actual disaster to occur, the customer should stress the importance of getting formal confirmation of the business continuity plan/disaster recovery (BCP/DR) tests, and whether the tests satisfied the SLAs contractually committed.

A customer should do the following:

- Review the contract of third-party commitments to maintain continuity of the provisioned service. The customer should also strongly consider further analysis.
- Review third-party business continuity processes and certifications.
- Coordinate a joint exercise with the CSP to test the customer's BCP/DR plan.
- Ensure that it receives confirmation of any BCP/DR tests the CSP undertakes, including results and lessons learned.
- Conduct an on-site assessment of the CSP facility to confirm and verify the asserted controls it uses to maintain the continuity of the service.

A CSP should do the following:

- Implement fast SLA-based data recovery. IaaS providers should have contractual agreements with multiple platform providers and have the tools in place to rapidly restore systems in the event of loss.
- Ensure incremental backups to frequently update a replica of all protected systems or snapshots at intervals set by the user for each system, so the consumer can determine the settings according to recovery point objectives.
- Make full site, system, disk and file recovery accessible via a user-driven, self-service portal that allows the user the flexibility to choose which file disk or system to recover.
- Negotiate the SLA up front and ensure that the customer pays for the SLA required to avoid conflict of interest.
- Consider adopting the most stringent requirements of any customer as a security baseline to maintain systems, facilities and procedures at a high level.

The controls in the CCM Business Continuity Management and Operational Resilience domain are based on the business continuity section of domain 6 in the CSA Guidance V4: Management Plane and Business Continuity.

Following are the controls:

- **BCR 01**—Business continuity planning
- **BCR 02**—Business continuity testing

[190] Singh, M.; "Gmail, Google Drive hit by outage," TechCrunch, 19 August 2020, https://techcrunch.com/2020/08/19/gmail-google-drive-down-users-say/

[191] Deoras, S.; "8 Cloud Outages That Shook The Tech World In 2019," Analytics India Magazine, 12 December 2019, https://analyticsindiamag.com/8-cloud-outages-that-shook-the-tech-world-in-2019/

- **BCR 03**—Data center utilities/environmental conditions
- **BCR 04**—Documentation
- **BCR 05**—Environmental risks
- **BCR 06**—Equipment location
- **BCR 07**—Equipment maintenance
- **BCR 08**—Equipment power failures
- **BCR 09**—Impact analysis
- **BCR 10**—Policy
- **BCR 11**—Retention policy

3.4.5 Change Control and Configuration Management

Configuration change controls for organizational information systems involve the systematic proposal, justification, implementation, testing, review and disposition of changes to the systems, including system upgrades and modifications.[192]

Configuration management is a collection of activities focused on establishing and maintaining the integrity of products and systems through control of the processes for initializing, changing and monitoring their configurations throughout the system development life cycle.[193]

A key component of a cloud security program is a rigorous change and configuration management system, including policies, procedures and technologies for managing the risks associated with applying changes to business-critical or customer/tenant-impacting assets (including applications, systems, infrastructure and configuration).

When considering change and configuration management, following are some key questions that any cloud actor (CSP, customer or auditor) should keep in mind:

- Who can request changes?
- Who can approve changes?
- Who can develop changes?
- Who can test the changes for compliance with approved specifications?
- Who can move the changes into production?

Some important requirements follow:

- Restrict programmer access to the production software and data to emergency basis only. This recommendation applies to developers for both CSP and CSC.
- Ensure that programmers cannot make uncontrolled changes to the source code's production version (see the DevOps section in chapter 8 in this study guide for further elaboration on this topic). Auditors need to check whether separate test environments exist so that development, testing, quality assurance assessments, staging environments and production activities are segregated by their functions. Auditors need to review the following:
  - If the source or object code is synchronized with the production code
  - If there is a production installation plan with an uninstallation contingency plan
- The organization (and the auditors) should evaluate if the CSPs use configuration management tools to enable the following:
  - Storing information about versions and builds of the software and testware

[192] NIST, "National Vulnerability Database," https://nvd.nist.gov/800-53/Rev4/control/CM-3
[193] NIST Computer Security Resource Center, "Glossary," https://csrc.nist.gov/glossary/term/configuration_management

- Traceability between software and testware and different versions or variants
- Tracking which versions belong with which configurations (e.g., operating systems, libraries, browsers)
- Building and releasing management
- Baselining (e.g., all the configuration items that make up a specific release)
- Access control (checking in and out).

In the Guidance V4, there is no dedicated domain on change control and configuration management. The controls of the CCM Change Control and Configuration Management domain refer to guidelines included in several of the Guidance's domains, primarily in domain 6 (Management Plane and Business Continuity), 7 (Infrastructure Security), 8 (Virtualization and Containers) and 10 (Application Security).

Following are the controls:

- **CCC 01**—New development/acquisition
- **CCC 02**—Outsourced development
- **CCC 03**—Quality testing
- **CCC 04**—Unauthorized software installations
- **CCC 05**—Production changes

3.4.6 Data Security and Information Life Cycle Management

As defined in the Guidance V4, the data security life cycle includes six phases from creation to destruction:

1. **Create**—Generation of new digital content, or the alteration/update/modification of existing content
2. **Store**—Commission of digital data to a storage repository (in most cases nearly simultaneous with creation)
3. **Use**—Viewing, processing or otherwise using data in some sort of activity, not including modification
4. **Share**—Making information accessible to others, such as between users, to customers or to partners
5. **Archive**—Taking data out of active use and placing it in long-term storage
6. **Destroy**—Permanently destroying data using physical or digital means (e.g., crypto shredding)

The relevance of the cycle when defining security control objectives is that it offers guidance on how to govern data during each step of accessing, processing and storing. Such a structured view and approach to data management is of critical importance since data breaches in their various forms (e.g., loss of archives, unauthorized access, malicious disclosure, unsecure deletion, processing not in compliance with applicable laws and regulations) are considered the number one risk in cloud computing.[194]

The CCM Data Security and Information Lifecycle Management domain aligns with the recommendation and guidelines included in domain 5 (Information Governance) of the Guidance.

Following are the controls:

- **DSI 01**—Classification
- **DSI 02**—Data inventory/flows
- **DSI 03**—eCommerce transactions
- **DSI 04**—Handling/labeling/security policy
- **DSI 05**—Nonproduction data

[194] Cloud Security Alliance, "Top Threats to Cloud Computing: Egregious Eleven," 6 August 2019, https://cloudsecurityalliance.org/artifacts/top-threats-to-cloud-computing-egregious-eleven/

- **DSI 06**—Ownership/stewardship
- **DSI 07**—Secure disposal

3.4.7 Data Center Security

The security of a modern data center must cover physical security, network security, and data and user security. In addition, there are physical security requirements related to offices and working facilities in general.

While physical security is an area often considered only marginally related to cloud computing, since it mainly pertains to the responsibilities of IaaS and hosting providers, it is important for cloud customers to do the following to avoid unpleasant surprises:

- Apply the necessary due diligence to verify that the provider applies adequate physical security measures.
- Take responsibility for the physical security in the working environment under the customer's direct control.

In general, when it comes to data center and physical security, particular emphasis should be given to asset classification and management, access control systems, physical and remote access, identification of the equipment, the possibility of using location-aware technologies, ingress and egress policies, and access logging.

The CCM data center security domain covers physical security controls and includes the following measures:

- Classifying tangible assets in terms of business criticality and data sensitivity
- Implementing physical security perimeters to safeguard personnel, data and information systems
- Policies and procedures for the relocation or transfer of hardware, software or data/information to an offsite or alternate location
- The use of location-aware technologies to validate connection authentication integrity
- Policies and procedures for the secure disposal of equipment used outside the organization's premises
- Policies and procedures for maintaining a safe and secure working environment in offices, rooms, facilities and secure areas storing sensitive information
- Control and monitoring of ingress and egress points of entry to prevent unauthorized access to facilities

It should be noted that the controls included in the CCM Data Center Security domain do not have a dedicated corresponding domain in the Guidance.

Following are the controls:

- **DCS 01**—Asset management
- **DCS 02**—Controlled access
- **DCS 03**—Equipment identification
- **DCS 04**—Off-site authorization
- **DCS 05**—Off-site equipment
- **DCS 06**—Policy
- **DCS 07**—Secure area authorization
- **DCS 08**—Unauthorized persons entry
- **DCS 09**—User access

3.4.8 Encryption and Key Management

Encryption is the conversion of plaintext to cipher text using a cryptographic algorithm.[195]

Key management includes activities involving the handling of cryptographic keys and other related security parameters (e.g., initialization vectors and passwords) during the entire life cycle of the keys, including their generation, storage, establishment, entry and output, and zeroization.[196]

Cryptography is an essential resource when it comes to protecting data and information.

It mitigates the risk of unauthorized access to confidential or regulated data (for instance, personal data) and ensures control over data integrity and repudiability. The development, implementation and maintenance of encryption and key management policies, procedures and techniques are critical to secure data in transit, in use and at rest. This is true in any data processing, storing and transferring environment, including the cloud.

Encryption options vary in cloud computing, largely depending on the specific service model, provider, and application/deployment model.

Customers should ask two essential questions:

- Where is the encryption taking place?
  - **For IaaS**—Is it on the client side? Server side? Proxy?
  - **For PaaS**—Is it on the application layer? Database? etc.
  - **For SaaS**—Is it provider managed? Proxy managed?
- Who is managing the keys?
  - CSP
  - Customer
  - Third party

Depending on the answers to these questions, the set of technical controls to be implemented can vary largely. Nevertheless, there are some key controls that should be in place regardless of the specific environment and implementation scenarios, including the following:

- Defining and implementing encryption and key management roles and responsibilities
- Using cryptographic libraries certified to approved standards (e.g., FIPS 140-3)
- Applying encryption based on data classification and risk
- Following standard change-management procedures
- Implementing, logging and monitoring the key lifecycle management
- Auditing encryption and key management systems and policies
- Implementing procedure for keys rotation and revocation

The CCM Encryption and Key Management domain aligns with the recommendations and guidelines included in domain 11 (Data Security and Encryption) of the Guidance.

Following are the controls:

- **EKM 01**—Entitlement
- **EKM 02**—Key generation

[195] NIST Computer Security Resource Center, "Glossary," https://csrc.nist.gov/glossary/term/encryption
[196] NIST Computer Security Resource Center, "Glossary," https://csrc.nist.gov/glossary/term/key_management

- **EKM 03**—Sensitive data protection
- **EKM 04**—Storage and access

3.4.9 Human Resources Security

The human factor is a key element in the development, implementation, provisioning and management of technology. Human resource (HR) management has to do with the rules, practices and actions that aid employees in effectively and efficiently participating in enterprise activities. Human resource management is about establishing, documenting, communicating, implementing, enforcing and maintaining policies, procedures and tools.

Human resource security is a key area in cybersecurity management, and in cloud security in particular. It has consistently been a factor in data breach reports[197, 198] and in the CSA Top Threats Reports.[199] Risk associated with the human factor relate to both malicious actions and employee errors due to lack of skill, expertise or awareness.

Following are examples of the measures the HR department should implement:

- Introduce provisions for adherence to established information governance and security policies.
- Assign, document and communicate the roles and responsibilities of employees, contractors and third-party users (e.g., through a RACI matrix).
- Create procedures to determine how and when enterprise-owned assets are to be returned upon contract termination.
- Introduce policies, procedures and technical measures to manage risks associated with permitting mobile device access to corporate resources (e.g., bring your own device (BYOD) policy).
- Enforce nondisclosure/confidentiality agreements.
- Establish and execute a security awareness training program for all contractors, third-party users and employees.
- Make personnel aware of their roles and responsibilities for maintaining awareness and compliance with established policies and procedures, and applicable legal, statutory or regulatory compliance obligations.

These controls apply both from the perspective of the CSP and the cloud customer.

The controls included in the CCM Human Resources Security domain do not have a direct corresponding domain in the Guidance.

Following are the controls:

- **HRS 01**—Asset returns
- **HRS 02**—Background screening
- **HRS 03**—Employment agreements
- **HRS 04**—Employment termination
- **HRS 05**—Mobile device management
- **HRS 06**—Nondisclosure Agreements
- **HRS 07**—Roles/responsibilities
- **HRS 08**—Technology acceptable use
- **HRS 09**—Training/awareness

[198] Verizon, “2018 Data Breach Investigations Report,” https://enterprise.verizon.com/resources/reports/DBIR_2018_Report.pdf
[199] *Op cit* Cloud Security Alliance, “Top Threats to Cloud Computing”

- **HRS 10**—User responsibility
- **HRS 11**—Workspace

3.4.10 Identity and Access Management

Cloud computing heavily impacts identity and access management as a security domain, mainly due to the increased complexity and increased exposure of the cloud computing model.

The increased complexity results from the cloud imposing shared responsibilities, collaboration and trusted relationships between multiple parties, i.e., between the IaaS, PaaS and SaaS providers, the end users and the identity providers. Such complex relationships are further complicated by the fact that very often an organization uses hundreds, if not thousands, of different services.

The increased exposure risk is related to the very nature of the cloud. The accessibility of services through the Internet dramatically increases the risks of credential theft, impersonification and account takeover, for example.

To properly implement IAM controls, an organization needs to map identities (such as a person, a device, an application, any other software, a virtual machine) to some attributes (e.g., the role within the company, the location of a server, the function of a piece of code). It is then necessary to define the right privileges associated with the identities.

The risks associated with IAM are that all of this takes place within a complex supply chain that includes several third parties that are not under the direct control of the organization, possibly using different systems and technologies. An identity mapping exercise would need to be coordinated within a defined service architecture that adopts identity federation by leveraging the right frameworks.

Other factors further complicating the IAM approach:

- Management of privileged users
- Management of third parties
- Employee-owned devices (due to BYOD policies)

Following are examples of measures an organization should have in place in order to mitigate the risks related to IAM:

- Establish, implement, enforce and maintain policies and procedures to manage information about the identities of who accesses data and enterprise-owned assets, and to determine their level of access.
- Review and update policies and procedures regularly.
- Enforce an authorization process based on the principles of least privilege and segregation of duties.
- Enforce strong authentication (e.g., multifactor authentication [MFA])
- Implement timely deprovisioning (revocation or modification) of access to data and enterprise-owned assets.
- Regularly review user entitlements.
- Ensure that identities are unique and traceable to an individual, credentials are security stored, and strong passwords are enforced.
- Restrict access to audit tools to prevent tampering or inappropriate disclosure of log data.
- Implement appropriate measures and compensating controls prior to provisioning third-party access.

The CCM Identity and Access Management domain aligns with the recommendations and guidelines included in domain 12 (Identity, Entitlement, and Access Management) of the Guidance.

Following are the controls:

- **IAM 01**—Audit tools access
- **IAM 02**—Credential lifecycle/provision management
- **IAM 03**—Diagnostic/configuration ports access
- **IAM 04**—Policies and procedures
- **IAM 05**—Segregation of duties
- **IAM 06**—Source code access restriction
- **IAM 07**—Third-party access
- **IAM 08**—Trusted sources
- **IAM 09**—User access authorization
- **IAM 10**—User access reviews
- **IAM 11**—User access revocation
- **IAM 12**—User ID credentials
- **IAM 13**—Utility programs access

3.4.11 Infrastructure and Virtualization

This CCM domain covers two fundamental pieces of cloud computing—the infrastructure and the virtualization technologies. The infrastructure is essentially the compute, storage and network physical layers, while the virtualization technologies are what allow the abstraction of the logical from the physical domain.

Virtualization is one of the key elements of IaaS cloud offerings and private clouds, and it is likely to be used in portions of the back end of PaaS and SaaS providers as well. Virtualization is, naturally, a key technology for virtual desktops, which are delivered from private or public clouds.

When addressing infrastructure and virtualization security in cloud computing, it is important to clearly understand the delineation of responsibilities based on the shared responsibility model and to determine what the CSP and customer should handle, respectively.

In general terms, the IaaS provider is in charge of the security of the virtualization infrastructure (e.g., physical security measures, network segregations, ensuring isolation, hardening the hypervisor).

The cloud customer is responsible for the security of the workloads (such as virtual machines or containers), the virtual network, images, security settings (e.g., managing access to the cloud management plane), logging and monitoring, and more. The cloud customer could be a PaaS or SaaS provider—i.e., any organization building upon the IaaS infrastructure.

Cloud infrastructures are rapidly evolving, and new concepts and technologies, such as serverless, are being adopted. Serverless technologies have redefined the attribution of the security responsibilities between IaaS and PaaS and SaaS, and between the CSP and customer.

Following are examples of measures an organization should have in place to mitigate the risk related to infrastructure and virtualization management.

- Establish, implement, enforce and maintain policies and procedures to ensure the security, retention and access control of audit logs.
- Continuously monitor security audit logs to detect anomalies, and take appropriate action.
- Monitor, encrypt and restrict communications between environments to authenticated and authorized connections.

- Document allowed services, protocols and ports.
- Harden the host and guest OS, hypervisor or infrastructure control plane.
- Design, develop, deploy and configure multitenant applications, infrastructure system and network components to ensure segmentation and segregation.
- Restrict access to network environments to authorized personnel.
- Provide segmentation, monitoring/detection and policy enforcement capabilities within the network architecture.
- Implement defense-in-depth techniques for detection and timely response to network-based attacks.

The CCM Infrastructure and Virtualization domain aligns with the recommendations and guidelines included in domains 7 and 8 (Infrastructure Security and Virtualization and Containers) of the Guidance.

Following are the controls:

- **IVS 01**—Intrusion detection
- **IVS 02**—Change detection
- **IVS 03**—Clock synchronization
- **IVS 04**—Information security documentation
- **IVS 05**—Vulnerability management
- **IVS 06**—Network security
- **IVS 07**—OS hardening and base controls
- **IVS 08**—Production/nonproduction environments
- **IVS 09**—Segmentation
- **IVS 10**—VM security—vMotion data protection
- **IVS 11**—VMM security—hypervisor hardening
- **IVS 12**—Wireless security
- **IVS 13**—Network architecture

3.4.12 Interoperability and Portability

Interoperability is the requirement that a processing system's components work together to achieve their intended result. It should be possible for the system to continue to work if components are replaced with new or different components from other providers.

Portability provides for application and data components to continue to work the same way when moved from one cloud environment to another without having to be changed. Portability is achieved by removing dependencies on the underlying environment. A portable component can be moved easily and reused regardless of the provider, platform, operating system, location, storage or other elements of the surrounding environment.

Portability and interoperability considerations are not unique to cloud environments, and their related security aspects are not new concepts resulting from cloud computing. However, the open and often shared processing environments that exist within the cloud require even greater precautions than traditional processing models. Multitenancy means that data and applications reside with data and applications of other companies, and that access to confidential data (intended or unintended) is possible through shared platforms, shared storage and shared networks (Guidance V3).

Lack of interoperability and portability can expose both cloud customers and CSPs to a variety of risks and considerably reduce the benefits of cloud computing. Among those risks:

- **Application, vendor or provider lock-in**—Choosing a particular cloud solution may restrict the customer's ability to move to another cloud offering or CSP.
- **Processing incompatibility and conflicts causing service disruption**—The provider, platform or application differences may expose incompatibilities that cause applications to malfunction within a different cloud infrastructure.
- **Unexpected application reengineering or business process changes**—Moving to a new cloud provider can introduce the need to rework the way a process functions or require coding changes to retain original behaviors.
- **Costly data migration or data conversion**—Lack of interoperable and portable formats may lead to unplanned data changes when a customer moves to a new provider.
- **Retraining or retooling**—New applications or management software may be required.
- **Loss of data or application security**—Differences in security policy, control, key management or data protection between providers may open undiscovered security gaps when a customer moves from one provider or platform to another.
- **Lack of compliance**—Moving to a new provider may result in conflicts with application laws and regulations.

Following are examples of measures that an organization should have in place to mitigate the risks related to the lack of interoperability and portability:

- Make published APIs securely available.
- Establish policies and procedures and define the terms of information processing interoperability and portability for application development.
- Implement standardized network protocols for the import and export of data.
- Implement standardized network protocols for the management of all interoperability and portability systems.
- Document relevant interoperability and portability standards in use.
- Use an industry-recognized virtualization platform and standard virtualization formats.

The controls included in the CCM Interoperability and Portability domain do not have a dedicated corresponding domain in the Guidance V4, but they do in V3 (domain 6, Interoperability and Portability).

Following are the controls:

- **IPY 01**—APIs
- **IPY 02**—Data request
- **IPY 03**—Policy and legal
- **IPY 04**—Standardized network protocols
- **IPY 05**—Virtualization

3.4.13 Mobile Security

Mobile computing is a broad term that can define any means of using a computer while outside the corporate office, including working from home, or from an airport or hotel while on the road.

The means to perform mobile computing can include kiosks to remotely connect to the corporate office, home computers, laptops, tablets or smartphones. Specialized or integrated devices could also be considered mobile computing devices.

The CCM Mobile Security domain focuses on implementing controls to mitigate the risks associated with module devices and end-point devices in general. Internet of Things (IoT) technologies are not in the scope of this domain.

The risk with mobile and end-point security mainly relate to user behavior, and the awareness (or lack) of a company's approach to acceptable use of devices and technologies (e.g., managed vs. unmanaged, enterprise-owned vs. personal).

Following are examples of measures an organization should take to mitigate the risks related to mobile and end-point devices:

- Maintain an inventory of all managed and unmanaged endpoints used to store and access company data.
- Deploy centralized solutions to manage and technically enforce policies and controls for all endpoints.
- Manage changes to endpoint operating systems and patch levels through the organization's change-management processes.
- Encrypt endpoints that are permitted to store company information to prevent unauthorized access to data in case the endpoint is lost, stolen or disposed of improperly.
- Establish, implement, enforce and maintain policies and procedures to define:
  - The acceptable use policy requirements, differentiating between managed and unmanaged endpoints
  - A list of the approved systems, servers, applications, application stores, application extensions and plugins
  - Prohibition of the installation of nonapproved applications
  - Prohibition of the circumvention of vendor-supported and integrated (built-in) security controls on endpoints (i.e., jailbreaking or rooting)
  - Privacy requirements for remote location identification, litigation, e-discovery, and legal holds, especially for personally owned devices.

The controls included in the CCM Mobile Security domain do not have a dedicated corresponding domain in the Guidance.

Following are the controls:

- **MOS 01**—Anti-malware
- **MOS 02**—Application stores
- **MOS 03**—Approved applications
- **MOS 04**—Approved software for BYOD
- **MOS 05**—Awareness and training
- **MOS 06**—Cloud-based services
- **MOS 07**—Compatibility
- **MOS 08**—Device eligibility
- **MOS 09**—Device inventory
- **MOS 10**—Device management
- **MOS 11**—Encryption
- **MOS 12**—Jailbreaking and rooting
- **MOS 13**—Legal
- **MOS 14**—Lockout
- **MOS 15**—Operating systems
- **MOS 16**—Passwords
- **MOS 17**—Policy
- **MOS 18**—Remote wipe

- **MOS 19**—Security patches
- **MOS 20**—Users

3.4.14 Security Incident Management, e-discovery and Cloud Forensics

An incident response plan is necessary for rapidly detecting incidents, minimizing loss and destruction, mitigating weaknesses exploited and restoring IT services.[200]

Numerous best practices and standards exist for incident management (handling/response), such as NIST 800-61rev2, ISO/IEC 27035 and the ENISA strategies for incident response and cybercrisis cooperation. The CCM takes these into account and aims to provide additional support specific to cloud computing.

The cloud computing environment imposes some changes compared to implementing an incident response plan in an on-premises scenario. This is mainly due to the need to determine and coordinate the actions of the CSP and the cloud customers. From a cloud customer perspective, incident response might become challenging if contracts and SLAs do not sufficiently clarify the roles and responsibilities of the parties, or if the communication and escalation processes (which are key in an effective incident response plan) lack specific direction on whom to contact and how in the event of a security event or incident.

Cloud forensics is the application of science to the identification, examination, collection and analysis of data while preserving the information and maintaining a strict chain of custody for the data in cloud computing (as a subset of network forensics).

Electronic discovery (e-discovery or eDiscovery) refers to any process in which electronic data is sought, located, secured and searched with the intent of using it as evidence in a civil or criminal legal case.

Forensics has to be carried out in a highly dynamic environment, which challenges basic forensic necessities such as establishing the scope of an incident, the collection and attribution of data, preserving the semantic integrity of the data, and maintaining the stability of evidence overall. These problems are exacerbated when cloud customers attempt to carry out forensic activities, since they operate in a nontransparent environment (which underscores the necessity of CSP support).

Following are examples of measures an organization should have in place to mitigate the risks related to incident management, e-discovery and forensics:

- Establish, document, implement, enforce and maintain security incident-response plans. Incident-response plans must be subject to testing for effectiveness at planned intervals or upon significant organizational or environmental changes.
- Ensure that incident response plans involve relevant internal departments, impacted customers, and other business relationships that represent critical intra-supply chain business process dependencies.
- Maintain points of contact for applicable regulation authorities, law enforcement and other legal jurisdictional authorities to ensure that direct compliance liaisons have been established and to be prepared for a forensic investigation.
- Establish, implement, enforce and maintain policies and procedures and supporting business processes and technical measures to triage security-related events and ensure timely and thorough incident management.
- Inform personnel and third parties of their responsibilities, and instruct them to report all information security events in a timely manner and through predefined communications channels.
- Establish, implement, enforce and maintain forensic procedures, including chain of custody.

[200] Cichonski, P.; T. Millar; T. Grance; K. Scarfone; “Computer Security Incident Handling Guide,” National Institute of Standards and Technology, August 2012, https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r2.pdf

- Implement mechanisms to monitor and quantify the types, volumes and costs of information security incidents.

The CCM Security Incident Management, E-Discovery and Cloud Forensics domain aligns with the recommendation and guidelines included in domains 3 and 9 (Legal Issues, Contracts and Electronic Discovery and Incident Response) of the Guidance.

Following are the controls:

- **SEF 01**—Contact/authority maintenance
- **SEF 02**—Incident management
- **SEF 03**—Incident reporting
- **SEF 04**—Incident response legal preparation
- **SEF 05**—Incident response metrics

3.4.15 Supply Chain Management, Transparency and Accountability

Supply chain management, transparency and accountability address the need for ensuring a CSP takes due care in managing its supply chain and the risks associated with governing data within the cloud. This domain is critical, because it deals with some of the core issues of cloud computing, such as the nontransferability of accountability, the need for transparency, and the impact of the cloud on security assessment, penetration testing and auditing.

Supply chain agreements (e.g., SLAs) between providers and customers should incorporate provisions or terms for the following:

- Scope of business relationship and services offered (e.g., customer data acquisition, exchange and usage, features and functionality, roles and responsibilities of provider and customer, subcontracted or outsourced business relationships, physical geographical location of hosted services, and any known regulatory compliance considerations)
- Primary points of contact for the duration of the business relationship
- References to detailed supporting and relevant business processes and technical measures implemented to enable effective governance, risk management, assurance, and legal, statutory and regulatory compliance obligations associated with all affected business relationships
- Notification and requisite preauthorization of any changes controlled by the provider with customer (tenant) impacts
- Timely notification of a security incident (or confirmed breach) to all customers (tenants) and other business relationships impacted (i.e., upstream and downstream impacted supply chain)
- Assessment and independent verification of compliance with agreement provisions and terms (e.g., industry-acceptable certification, attestation audit report or equivalent forms of assurance) without posing an unacceptable business risk of exposure to the organization being assessed
- Expiration of the business relationship and treatment of customer (tenant) data impacted
- Customer service-to-service application (API) and data interoperability and portability requirements for application development and information exchange, usage and integrity persistence

It is critical that CSPs clearly delineate shared responsibility model control applicability and ownership for each control for their specific service model and offering, including which controls have customer security responsibilities (in whole or in part). They should clearly and completely describe those responsibilities.

For the customers it is important to review and validate the CSP's shared responsibility model documentation. Both parties must implement and operate their instances of the service accordingly.

The controls included in the CCM Supply Chain Management, Transparency and Accountability domain do not have a dedicated corresponding domain in the Guidance.

Following are the controls:

- **STA 01**—Data quality and integrity
- **STA 02**—Incident reporting
- **STA 03**—Network/infrastructure services
- **STA 04**—Provider internal assessments
- **STA 05**—Supply chain agreements
- **STA 06**—Supply chain governance reviews
- **STA 07**—Supply chain metrics
- **STA 08**—Third-party assessment
- **STA 09**—Third-party audits

3.4.16 Threat and Vulnerability Management

Threat and vulnerability management programs provide a way to assess the potential business impact and likelihood of threats and risk to the organization information infrastructure before events occur.[201]

Following are examples of measures that an organization should include in its threat and vulnerabilities management program:

- Establish, implement, enforce and maintain policies, procedures and technical measures to prevent the execution of malware on enterprise-owned or managed assets.
- Implement technical measures and supporting processes for vulnerabilities detection and use a risk-based model for prioritizing the remediation of identified vulnerabilities.
- Implement a process for tracking and reporting vulnerability identification and remediation activities.
- Provide a summary of identified weaknesses to stakeholders if the system owner shares responsibility for the remediation, and inform them of the policies and procedures for threat and vulnerability management.
- Update detection tools, threat signatures and indicators of compromise.
- Perform periodic penetration testing.
- Establish metrics for vulnerability identification and remediation at defined intervals.

The controls included in the CCM Threat and Vulnerability Management domain do not have a dedicated corresponding domain in the Guidance.

Following are the controls:

- **TVM 01**—Antivirus/malicious software
- **TVM 02**—Vulnerability/patch management
- **TVM 03**—Mobile code

[201] Dildy, T.; "Enterprise Vulnerability Management," *ISACA® Journal*, 31 March 2017, www.isaca.org/resources/isaca-journal/issues/2017/volume-2/enterprise-vulnerability-management

3.4.17 Governance and Risk Management

Corporate governance is the set of processes, technologies, customs, policies, laws and institutions affecting the way an enterprise is directed, administered or controlled. It also includes the relationship among the many stakeholders involved and the goals of the company.

Information risk management is the process of identifying and understanding exposure to risk and the capability of managing it, aligned with the data owner's risk appetite and tolerance.

An effective governance and enterprise risk management cloud computing program flows from well-developed information security governance processes as part of the organization overall corporate governance obligations of due care.

This CCM domain aligns with the recommendation and guidelines included in domain 2 (Governance and Enterprise Risk Management) of the Guidance.

See chapter 1 on cloud governance for more details.

Following are the controls:

- **GRM 01**—Baseline Requirements
- **GRM 02**—Data focus risk assessments
- **GRM 03**—Management oversight
- **GRM 04**—Management program
- **GRM 05**—Management support/involvement
- **GRM 06**—Policy
- **GRM 07**—Policy enforcement
- **GRM 08**—Policy impact on risk assessments
- **GRM 09**—Policy reviews
- **GRM 10**—Risk assessments
- **GRM 11**—Risk management framework

3.5 The Consensus Assessment Initiative Questionnaire (CAIQ)

This section describes the Consensus Assessment Initiative Questionnaire (CAIQ).

3.5.1 The CAIQ and Why It Was Created

The Consensus Assessment Initiative Questionnaire (CAIQ) provides cloud customers and auditors with a set of questions they can ask a CSP about its security posture and adherence to the CSA best practices (CCM and the Guidance). The CAIQ was created as a companion document to support better adoption of the CCM. Where the CCM defines the control specification, the CAIQ defines questions to ensure implementation.

The CAIQ consists of yes or no answers that can be expanded on with explanatory text in a dedicated notes column. This allows a CSP to provide additional information that a potential customer can use to better evaluate how a control has been implemented. It also enables the provider to eventually offer a better explanation of the application of the shared responsibility model. Like the CCM, the CAIQ comes in a spreadsheet format with a structure that is identical to the Cloud Controls Matrix.

The relationship between a CCM control and a CAIQ question in most cases is many to one. This is by design, because the CCM is based on 133 controls, whereas the CAIQ has 310 questions in its latest version 3.1. Depending on the nature and the complexity of the CCM control, there can be one or several questions to ask to verify the implementation of a certain control.

For example, the control AAC-02 in the Audit and Assurance Compliance domain (see **figure 3.3**), includes numerous requirements. Therefore, more than one question is required to verify whether the CSP has implemented all the necessary requirements within that control.

Figure 3.3—Control AAC-02 Example

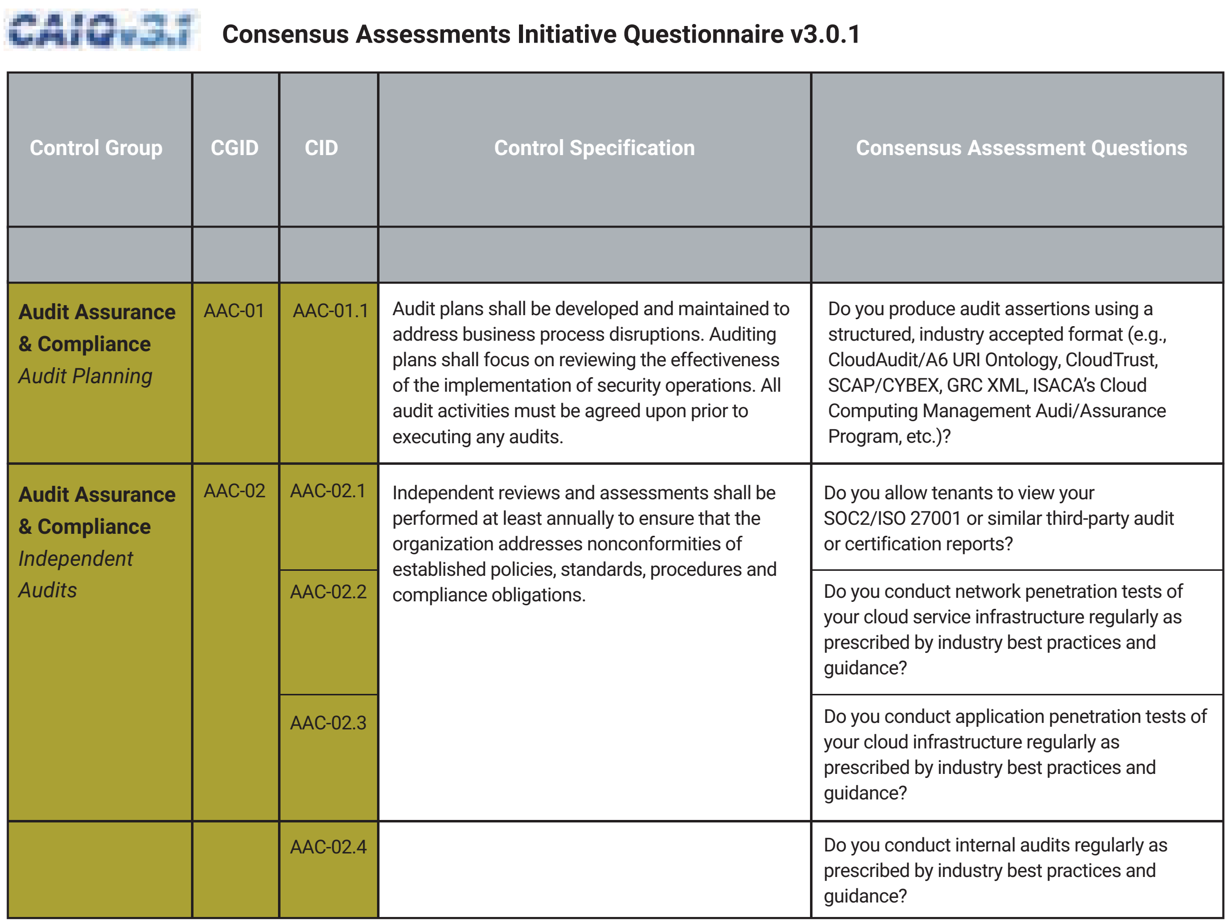

Consensus Assessments Initiative Questionnaire v3.0.1

Control Group	CGID	CID	Control Specification	Consensus Assessment Questions
Audit Assurance & Compliance *Audit Planning*	AAC-01	AAC-01.1	Audit plans shall be developed and maintained to address business process disruptions. Auditing plans shall focus on reviewing the effectiveness of the implementation of security operations. All audit activities must be agreed upon prior to executing any audits.	Do you produce audit assertions using a structured, industry accepted format (e.g., CloudAudit/A6 URI Ontology, CloudTrust, SCAP/CYBEX, GRC XML, ISACA's Cloud Computing Management Audi/Assurance Program, etc.)?
Audit Assurance & Compliance *Independent Audits*	AAC-02	AAC-02.1	Independent reviews and assessments shall be performed at least annually to ensure that the organization addresses nonconformities of established policies, standards, procedures and compliance obligations.	Do you allow tenants to view your SOC2/ISO 27001 or similar third-party audit or certification reports?
		AAC-02.2		Do you conduct network penetration tests of your cloud service infrastructure regularly as prescribed by industry best practices and guidance?
		AAC-02.3		Do you conduct application penetration tests of your cloud infrastructure regularly as prescribed by industry best practices and guidance?
		AAC-02.4		Do you conduct internal audits regularly as prescribed by industry best practices and guidance?

Source: Cloud Security Alliance, Cloud Controls Matrix (CCM) Working Group

3.5.2 CAIQ Target Audience, Purpose and Differences From CCM

This section describes the CAIQ target audience, purpose and differences from CCM.

Target Audience

The CAIQ target audience includes cloud customers, CSPs and auditors.

CAIQ for Cloud Customers

Customers can use the standard set of questions in the CAIQ to verify that a CSP meets their security specifications. Cloud customers can use the CAIQ as a tool throughout various stages of the cloud lifecycle.

During the pre-vetting of a CSP, before entering into an agreement, customers can use the CAIQ to compare provider security and privacy features with its own internal requirements. Once the service is up and running, customers can use the CAIQ to monitor if the security controls' objectives are met over time.

CAIQ for CSPs

CSPs use the CAIQ for two purposes: Generate and provide controls and risk models (like the CCM), and generate a template questionnaire that is distributed to cloud customers.

Service providers receive many customer requests for personalized questionnaires. This creates a significant workload, which is why CSPs started using CAIQ as a way to streamline their approach to customer requests. Now, instead of filling out hundreds (or even thousands) of customer requests, they get ahead by filling out one standard questionnaire and sending it to all their customers.

This same issue also spurred the creation of the STAR Registry. Providers can post the CAIQ to this registry to demonstrate the breadth of their control, and as a standard response to an RFP. Examples of these (such as Microsoft Office 365) can be found in the STAR registry.[202]

CAIQ for Auditors

From an auditor's standpoint, the use of CAIQ is straightforward—to pose the right questions to the auditee. For an internal auditor, this can serve as a guide for questions to ask during an assessment; for external auditors evaluating the service provider in the context of an internal evaluation or an audit for a certification or attestation, the CAIQ can serve as a supplement to the certification or attestation standard.

Filling Out the Questionnaire

CAIQ collects data from CSPs about the extent of their alignment with the CCM requirements, their compliance with regulations and frameworks, and the security of their infrastructure. The CAIQ questions are grouped into control areas that probe the state of implementation.

The CSPs are expected to answer questions as follows:

- Yes, if a control is implemented
- No, if a control is not implemented
- Not applicable, if a control is not in scope or relevant (e.g., a control on data center security might not be directly applicable to a SaaS provider, which is likely to be accountable only for its monitoring compliance inheritance, but not responsible for the actual implementation.)

In addition, the CAIQ template includes a free text box for a CSP to provide additional information potentially useful to a cloud customer during its evaluation process, such as documenting why a particular control is not applicable.

Figure 3.4 is an example from the CSA STAR Registry.

[202] Cloud Security Alliance, "CSA STAR Registry," https://cloudsecurityalliance.org/star/registry/

Figure 3.4—CSA STAR Registry Example

Audit Assurance & Compliance *Independent Audits*	AAC-02	AAC-02.1	Independent reviews and assessments shall be performed at least annually to ensure that the organization addresses nonconformities of established policies, standards, procedures, and compliance obligations.	Do you allow tenants to view your SOC2/ISO 27001 or similar third-party audit or certification reports?	Yes	ISO 27001 certifications for XYZ and CO+I can be found on the XYZ Trust Portal at https://servicetrust.XYZ.com/. Customers can also review SOC, ISO, PCI and other audit reports. Additional audit information is available under NDA upon request by prospective and existing customers through their XYZ account representative.
		AAC-02.2		Do you conduct network penetration tests of your cloud service infrastructure at least annually?	Yes	As defined in AIS-01.1, regular scans are conducted, at least quarterly, against the Azure infrastructure and applications using a variety of commercial and proprietary scanning tools. Critical and High findings detected are reviewed and patched per the Change and Release Management Policy. Re-scans are conducted within 30 days. Assume Breach employs Red-Team/Blue-Team exercises, live site penetration testing and centralized security logging and monitoring to identify and address potential gaps, test security response plans, reduce exposure to attack, and reduce access from a compromised system, with periodic post-breach assessment and clean state restoration.
		AAC-02.3		Do you conduct application penetration tests of your cloud infrastructure regularly as prescribed by industry best practices and guidance?	Yes	Scans are performed by Co+I security professionals on behalf of XYZ. Penetration testing methodologies for Infrastructure and Application are defined and are based on a combination of common criteria, NIST SP800-115, ETSI, OWASP, IETF and ISO 27000. To protect XYZ platform services, XYZ provides a distributed denial-of-service

Figure 3.4—CSA STAR Registry Example

						(DDoS) defense system that is part of XYZ's continuous monitoring process, and is continually improved through scheduled penetration-testing and red team exercises. The XYZ DDoS defense system is designed to mitigate attacks from the outside and also from other XYZ tenants. The XYZ DDoS defense technology provides detection and mitigation techniques such as SYN cookies, rate limiting and connection limits to help ensure that such attack do not impact the customer environment.

The example considers three questions from the AAC Domain—AAC-02.1, AAC-02.2 and AAC-02.3. The CSP offers several additional details that provide the customer better visibility into the CSP practices. It also provides the auditor with better data for the evidence collection process.

Using CAIQ, an organization can build a robust RFP (request for proposal) and verify that the answers the CSP gives during the RFP interview are valid. Even so, it is always a good idea to perform due diligence if required, as it is easy for a respondent to check boxes. Whether the CSP is following through with the implementation should be confirmed.

Differences Between CCM and CAIQ

The CCM and CAIQ help provide a solid foundation for assessing the risk models and controls of a cloud provider, but there are some differences between the two tools (**figure 3.5**).

Figure 3.5—Comparison of CCM and CAIQ

CCM™

- Provides a controls framework to give a detailed understanding of security concepts and principles that are aligned to the Guidance domains
- Assesses the overall security of a CSP service
- Covers 16 domains
- Includes 133 controls
- Specifies overall security needs of cloud consumers
- Assesses the overall security risk of a cloud provider

CAIQ™

- Provides industry-accepted ways to document which security controls exist in IaaS, PaaS and SaaS offerings
- Assesses presence of controls of a CSP's service
- Includes 310 questions
- Covers 133 controls
- A cloud consumer and cloud auditor may wish to review a CSP to gauge the CSP's security and conduct initial assessment followed by further clarifying questions

Source: Cloud Security Alliance, Cloud Controls Matrix (CCM) Working Group

3.6 CCM and CAIQ Structure

This section describes the CCM and CAIQ structure.

3.6.1 CCM Structure

The CCM V3.0.1 is structured with 16 security domains and 133 controls.

Control Domains

Domains are color coded in the spreadsheet for easy identification (**figure 3.6**). Under each domain (such as Application and Interface Security, Business Continuity, Encryption and Key Management), there are several controls that can be selected for implementation depending on the environment, or for auditing cloud providers.

Figure 3.6—CCM Security Domains

AIS	Application & Interface Security	HRS	Human Resources Security
AAC	Audit Assurance & Compliance	IAM	Identity & Access Management
BCR	Business Continuity Mgmt & Op Resilience	IVS	Infrastructure & Virtualization
CCC	Change Control & Configuration Management	IPY	Interoperability & Portability
DSI	Data Security & Information Lifecycle Mgmt	MOS	Mobile Security
DCS	Datacenter Security	SEF	Sec. Incident Mgmt, E-Disc & Cloud Forensics
EKM	Encryption & Key Management	STA	Supply Chain Mgmt, Transparency & Accountability
GRM	Governance & Risk Management	TVM	Threat & Vulnerability Management

Source: Cloud Security Alliance, Cloud Controls Matrix (CCM) Working Group

CCM Columns

Figure 3.7 shows the CCM controls matrix.

Figure 3.7—Cloud Controls Matrix

Cloud Controls Matrix																
Control Domain	CCM 3.0 Control ID	Updated Control Specification	Architectural Relevance						Corp Gov Relevance	Cloud Service Delivery Model Applicability			Supplier Relationship		Scope Applicability	
			Phys	Network	Computer	Storage	App	Data		SaaS	PaaS	IaaS	Service Provider	Tenant/ Consumer	ISO/IEC 27001 – 2013	NIS SP800-53 R4 App J
Application & Interface Security Customer Access Requirements	AIS-02	Prior to granting customers access to data, assets, and information systems, security, contractual, and regulatory requirements for customer access shall be addressed.	X	X	X	X	X	X	X	X	X	X	X	X	A9.1.1	AP-1 The organization determines and documents the legal authority that permits the collection, use, maintenance, and sharing of personal identifiable information (PII), either generally or in support of a specific program or information system need.

Source: Cloud Security Alliance, Cloud Controls Matrix (CCM) Working Group

Control Domain and ID Column

Figure 3.8 shows the CCM control domain and ID column.

Figure 3.8—Control Domain and ID Column Example

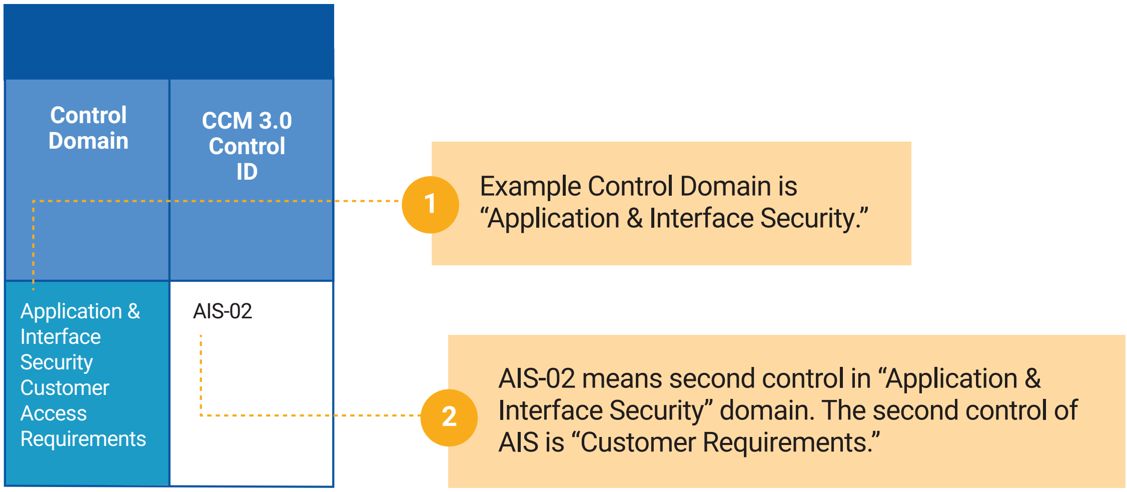

Control Domain	CCM 3.0 Control ID
Application & Interface Security Customer Access Requirements	AIS-02

1 Example Control Domain is "Application & Interface Security."

2 AIS-02 means second control in "Application & Interface Security" domain. The second control of AIS is "Customer Requirements."

Source: Cloud Security Alliance, Cloud Controls Matrix (CCM) Working Group

Control Domain

The domain defines which category the controls fall under. Each domain is assigned an acronym. For example, the Application and Interface Security domain is assigned the acronym AIS.

Control ID

There may be one or several controls under a domain. Each control has a control ID, which is the domain acronym followed by a number. For example, AIS-02 means second control in the Application and Interface Security domain. Another example is IAM-03. This means third control in the Identity and Access Management domain.

Control Specification and Architectural Relevance

The Control Specification column describes the purpose of the control (**figure 3.9**).

Figure 3.9—Control Specification and Architectural Relevance Example

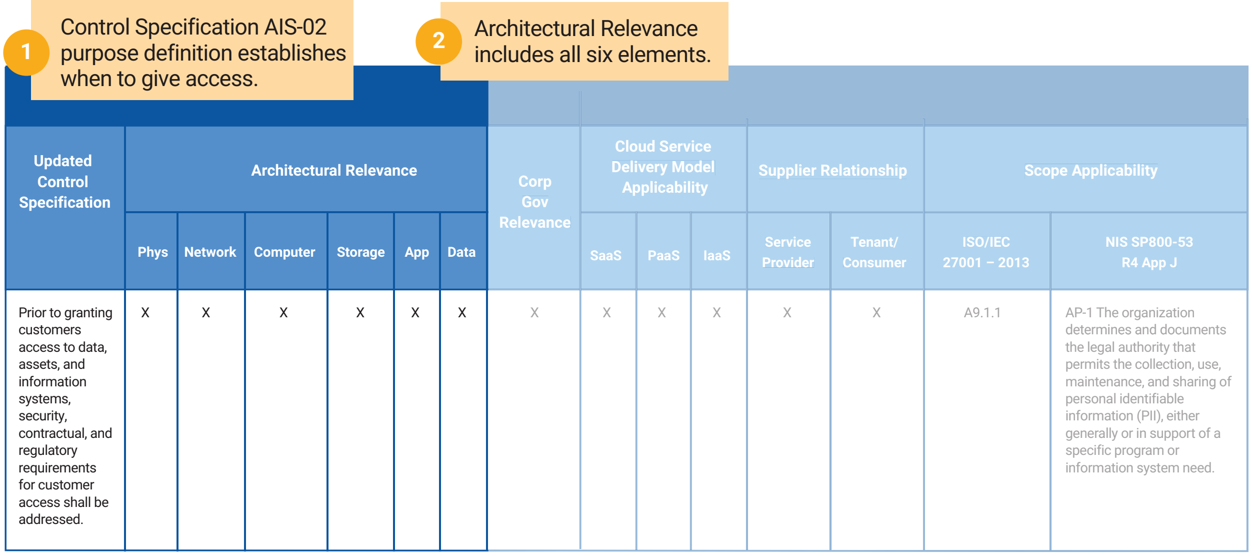

1 Control Specification AIS-02 purpose definition establishes when to give access.

2 Architectural Relevance includes all six elements.

Updated Control Specification	Architectural Relevance						Corp Gov Relevance	Cloud Service Delivery Model Applicability			Supplier Relationship		Scope Applicability	
	Phys	Network	Computer	Storage	App	Data		SaaS	PaaS	IaaS	Service Provider	Tenant/ Consumer	ISO/IEC 27001 – 2013	NIS SP800-53 R4 App J
Prior to granting customers access to data, assets, and information systems, security, contractual, and regulatory requirements for customer access shall be addressed.	X	X	X	X	X	X	X	X	X	X	X	X	A9.1.1	AP-1 The organization determines and documents the legal authority that permits the collection, use, maintenance, and sharing of personal identifiable information (PII), either generally or in support of a specific program or information system need.

Source: Cloud Security Alliance, Cloud Controls Matrix (CCM) Working Group

Architectural Relevance

The CCM has six columns for architectural relevance, which includes applicability to the physical infrastructure, network, compute, storage, applications, data, or all of these areas. Applicability of architectural relevance to the corresponding control is marked with an X in the column. Columns can be used for filtering (i.e., all storage elements).

Depending on the description in the control specification, the volunteers at CSA determined where a control is applicable. Does it apply to physical infrastructure? Applications? Data? Or does it apply to all these areas? Because customer access needs to be considered for all these areas, there is an X in all these columns.

Some controls apply only to specific areas. For example, security of ecommerce data that traverses public networks will need to be appropriately classified and protected from fraudulent activity, unauthorized disclosure, or modification in such a manner to prevent contract dispute and compromised data. This control is applicable only to network and data, as the other areas, like storage and physical infrastructure, are not directly affected by this control.

Corporate Governance Relevance

Corporate governance relevance (**figure 3.10**):

- Determines whether the control is relevant to corporate governance.
- Controls that are relevant to corporate governance are marked with an X.

Figure 3.10—Corporate Governance Relevance

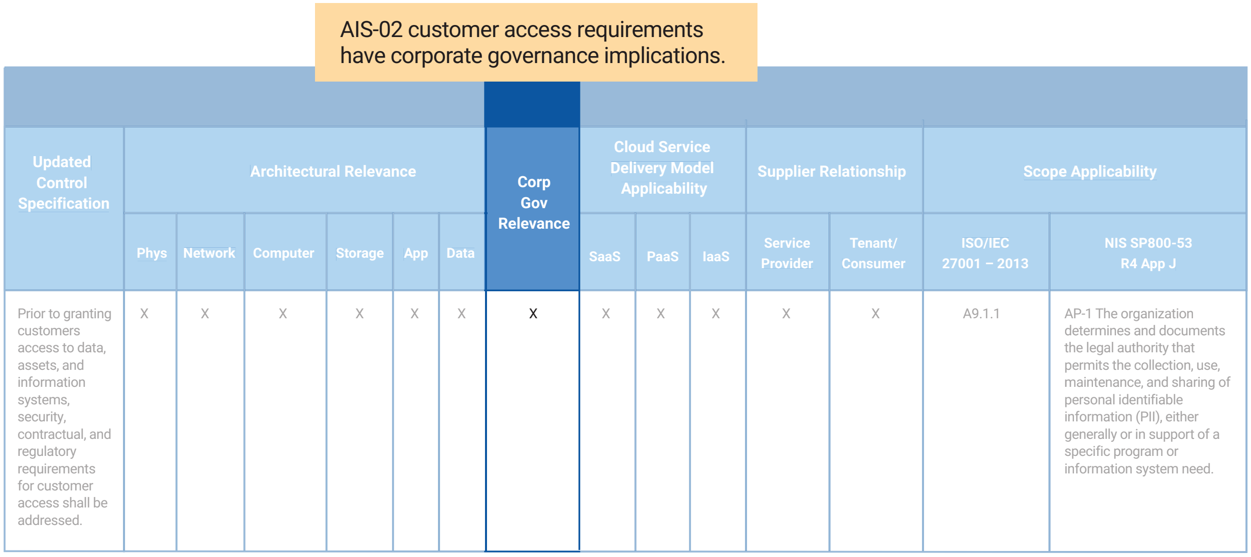

AIS-02 customer access requirements have corporate governance implications.

Updated Control Specification	Architectural Relevance						Corp Gov Relevance	Cloud Service Delivery Model Applicability			Supplier Relationship		Scope Applicability	
	Phys	Network	Computer	Storage	App	Data		SaaS	PaaS	IaaS	Service Provider	Tenant/ Consumer	ISO/IEC 27001 – 2013	NIS SP800-53 R4 App J
Prior to granting customers access to data, assets, and information systems, security, contractual, and regulatory requirements for customer access shall be addressed.	X	X	X	X	X	X	X	X	X	X	X	X	A9.1.1	AP-1 The organization determines and documents the legal authority that permits the collection, use, maintenance, and sharing of personal identifiable information (PII), either generally or in support of a specific program or information system need.

Source: Cloud Security Alliance, Cloud Controls Matrix (CCM) Working Group

As defined in chapter 1, corporate governance is the set of processes, customs, policies, laws and institutions affecting the way an enterprise is directed, administered or controlled. Corporate governance includes the relationship among the many stakeholders involved and the goals of the company. Good governance is based on the acceptance of the rights of shareholders, as the true owners of the corporation, and the role of senior management as trustees.

Cloud Service Delivery Model Applicability

Cloud service delivery model applicability (**figure 3.11**) determines which service model controls are applicable to IaaS, PaaS or SaaS**.**

Figure 3.11—Cloud Service Delivery Model Applicability

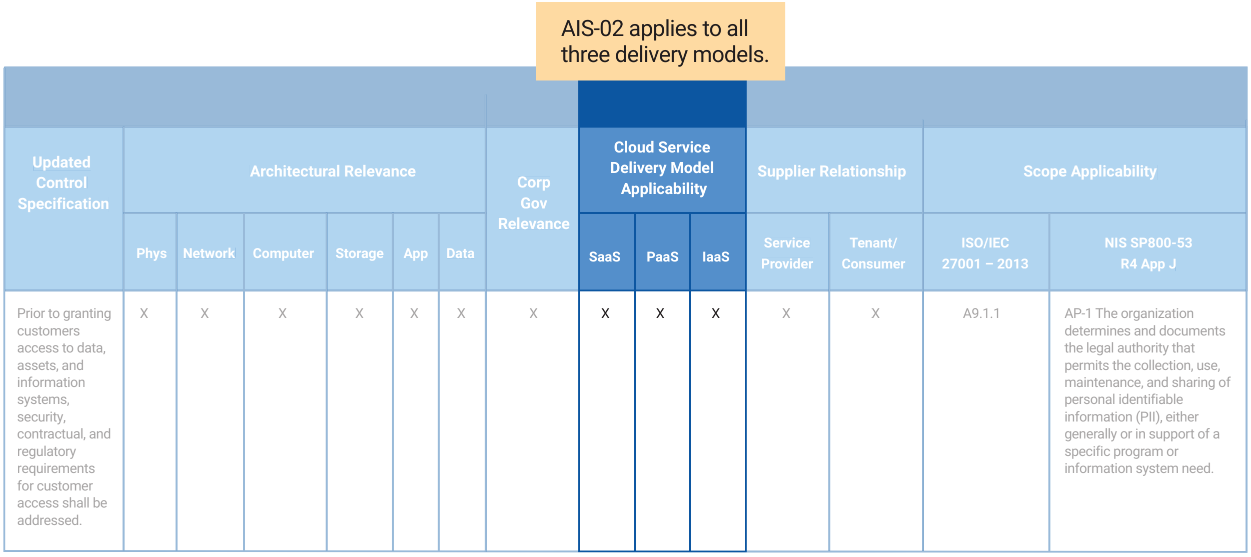

AIS-02 applies to all three delivery models.

Updated Control Specification	Architectural Relevance						Corp Gov Relevance	Cloud Service Delivery Model Applicability			Supplier Relationship		Scope Applicability	
	Phys	Network	Computer	Storage	App	Data		SaaS	PaaS	IaaS	Service Provider	Tenant/ Consumer	ISO/IEC 27001 – 2013	NIS SP800-53 R4 App J
Prior to granting customers access to data, assets, and information systems, security, contractual, and regulatory requirements for customer access shall be addressed.	X	X	X	X	X	X	X	X	X	X	X	X	A9.1.1	AP-1 The organization determines and documents the legal authority that permits the collection, use, maintenance, and sharing of personal identifiable information (PII), either generally or in support of a specific program or information system need.

Source: Cloud Security Alliance, Cloud Controls Matrix (CCM) Working Group

This column explains which model a control applies to—IaaS, PaaS or SaaS. For example, this would be applicable if a cloud customer wanted to assess a cloud provider to determine if it wanted to use its IaaS service. Microsoft, for example, has used the CCM to address security concerns for some of its products, such as Azure, Microsoft Dynamics 365, and Office 365. A customer seeking to assess an Azure deployment for its IaaS environment would look at the X mark in the IaaS column to identify which controls apply, and based on that could assess the security of Microsoft's IaaS services.

Supplier Relationship

The supplier relationship (**figure 3.12**):

- Determines the applicability of the controls to the provider or the tenant.
- Applicability of supplier relationship to the corresponding control is marked with an X in the column.

Figure 3.12—Supplier Relationship

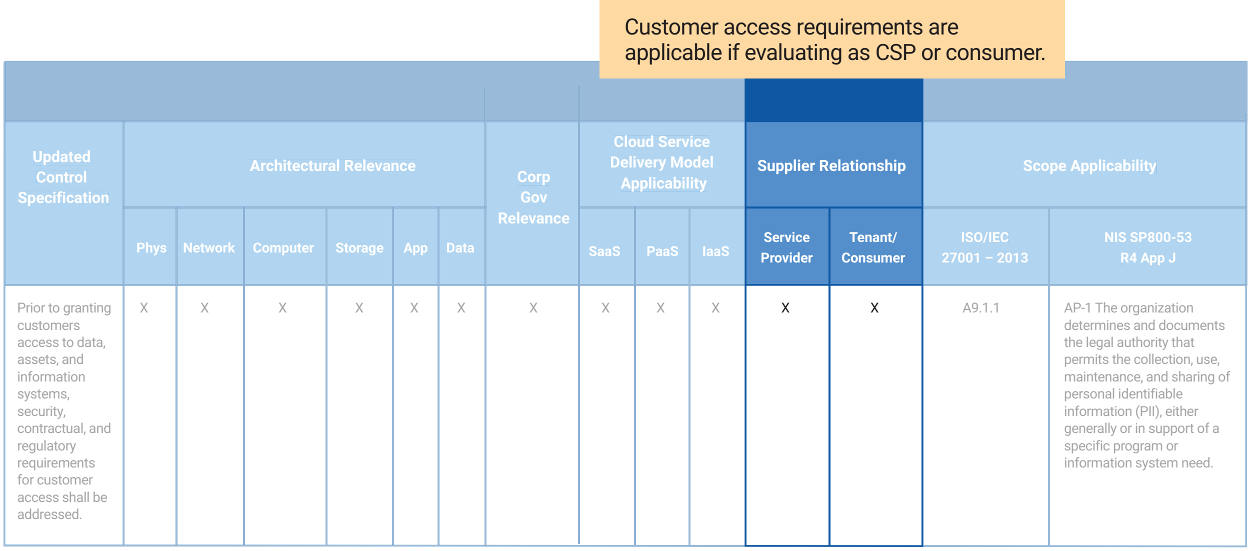

Customer access requirements are applicable if evaluating as CSP or consumer.

Updated Control Specification	Architectural Relevance						Corp Gov Relevance	Cloud Service Delivery Model Applicability			Supplier Relationship		Scope Applicability	
	Phys	Network	Computer	Storage	App	Data		SaaS	PaaS	IaaS	Service Provider	Tenant/ Consumer	ISO/IEC 27001 – 2013	NIS SP800-53 R4 App J
Prior to granting customers access to data, assets, and information systems, security, contractual, and regulatory requirements for customer access shall be addressed.	X	X	X	X	X	X	X	X	X	X	X	X	A9.1.1	AP-1 The organization determines and documents the legal authority that permits the collection, use, maintenance, and sharing of personal identifiable information (PII), either generally or in support of a specific program or information system need.

Source: Cloud Security Alliance, Cloud Controls Matrix (CCM) Working Group

This column determines the supplier relationship. For example, controls related to business continuity and management of the data center are applicable to the provider and not the tenant. There is an X in the supplier cell and the tenant cell will be blank. This column helps in the adoption of the shared responsibility model, because it makes clear who is responsible for which controls (in terms of design, implementation, testing and management).

Scope Applicability

Scope applicability (**figure 3.13**):

- Maps existing industry frameworks to controls in the 16 domains.
- Mappings leverage the work done with other standards, regulations and control frameworks.

Figure 3.13—Scope Applicability

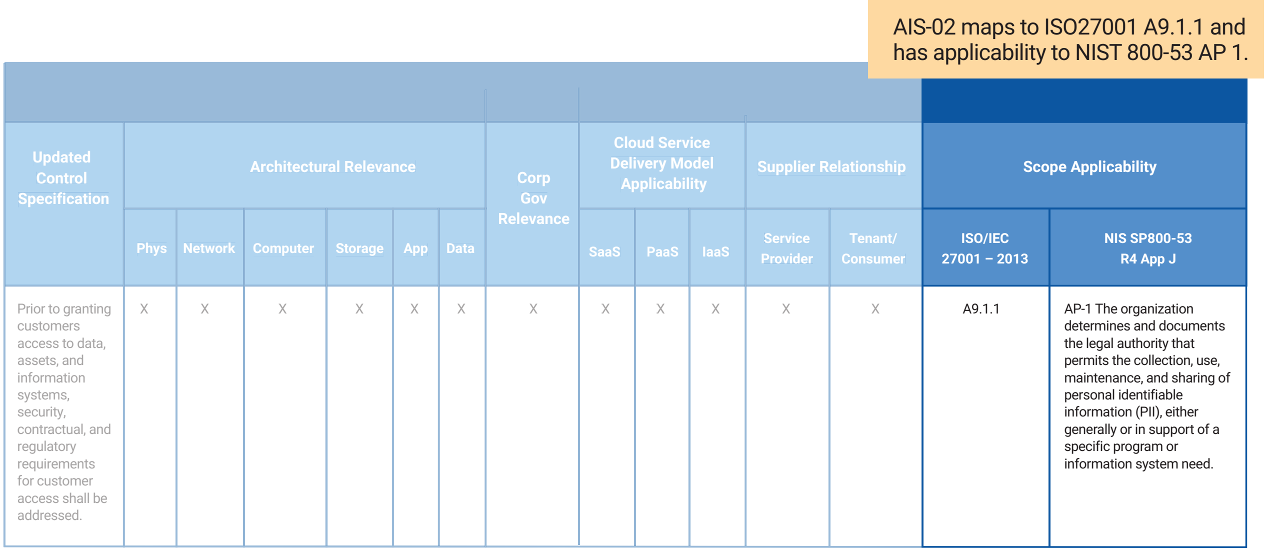

AIS-02 maps to ISO27001 A9.1.1 and has applicability to NIST 800-53 AP 1.

Updated Control Specification	Architectural Relevance						Corp Gov Relevance	Cloud Service Delivery Model Applicability			Supplier Relationship		Scope Applicability	
	Phys	Network	Computer	Storage	App	Data		SaaS	PaaS	IaaS	Service Provider	Tenant/ Consumer	ISO/IEC 27001 – 2013	NIS SP800-53 R4 App J
Prior to granting customers access to data, assets, and information systems, security, contractual, and regulatory requirements for customer access shall be addressed.	X	X	X	X	X	X	X	X	X	X	X	X	A9.1.1	AP-1 The organization determines and documents the legal authority that permits the collection, use, maintenance, and sharing of personal identifiable information (PII), either generally or in support of a specific program or information system need.

Source: Cloud Security Alliance, Cloud Controls Matrix (CCM) Working Group

The last column in the CCM maps the existing industry frameworks to the controls of the 16 domains. This is one of the most useful features of the CCM, as it leverages the work done by existing standards, regulations and control frameworks.

Currently, the CCM controls are mapped to approximately 40 standards. As the previous example illustrates, the ISO standard is mapped using the section number A9.1.1. Some standards have the actual control mentioned, in this case NIST. For example, if the cloud provider being assessed is compliant with the NIST SP800-53 standard, it probably already has measures in place to fulfill this requirement.

3.6.2 CAIQ Structure

Figure 3.14 shows the Consensus Assessments Initiative Questionnaire (CAIQ) structure.

Figure 3.14—CAIQ Structure

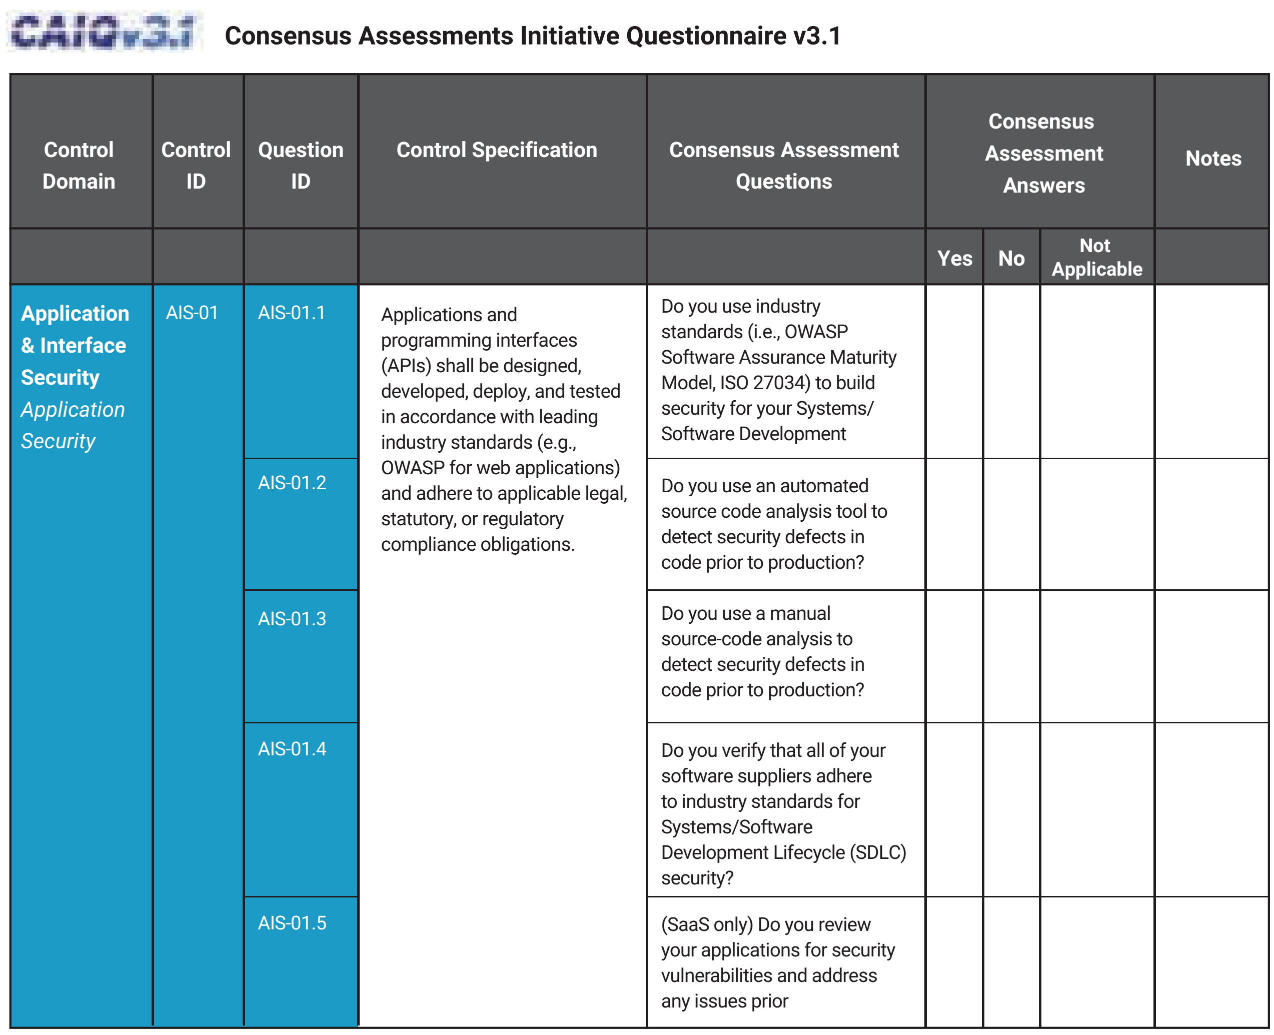

CAIQv3.1 **Consensus Assessments Initiative Questionnaire v3.1**

Control Domain	Control ID	Question ID	Control Specification	Consensus Assessment Questions	Consensus Assessment Answers			Notes
					Yes	No	Not Applicable	
Application & Interface Security *Application Security*	AIS-01	AIS-01.1	Applications and programming interfaces (APIs) shall be designed, developed, deploy, and tested in accordance with leading industry standards (e.g., OWASP for web applications) and adhere to applicable legal, statutory, or regulatory compliance obligations.	Do you use industry standards (i.e., OWASP Software Assurance Maturity Model, ISO 27034) to build security for your Systems/ Software Development				
		AIS-01.2		Do you use an automated source code analysis tool to detect security defects in code prior to production?				
		AIS-01.3		Do you use a manual source-code analysis to detect security defects in code prior to production?				
		AIS-01.4		Do you verify that all of your software suppliers adhere to industry standards for Systems/Software Development Lifecycle (SDLC) security?				
		AIS-01.5		(SaaS only) Do you review your applications for security vulnerabilities and address any issues prior				

Source: Cloud Security Alliance, Cloud Controls Matrix (CCM) Working Group

- **Control Domain**—Same as the Control Domain in CCM (color coded to match to CCM v3.0.1)
- **Control ID**—A unique ID that identifies each control within a group or domain (same as the control ID in CCM v3.0.1)
- **Question ID**—The question ID for that control
- **Control specification**—Describes the purpose of the control (same as the control specification in CCM v3.0.1)
- **Consensus assessment questions**—Questions corresponding to that control
- **Consensus assessment answers**—Three options: Yes, No and N/A, with explanations encouraged in each case
- **Notes**—For any additional information or comment

In some cases, a question may not be applicable to a cloud provider. For example, if a cloud provider offers only PaaS services and uses another provider for its infrastructure needs, some data center control questions may not be applicable.

3.7 CCM Relationship With Other Frameworks: Mappings and Gap Analysis

This section describes the CCM relationship with other frameworks.

3.7.1 CCM Mapping Objective and Scope

One of the key features of the CCM is the mapping with other frameworks. Although CCM is cloud-specific, it is normally integrated into organizations where other frameworks are being used. Because of this, it is paramount to establish a link between the standards in use and the CCM. The main objective of the mapping process is to map the CCM's controls to other frameworks and identify a semantic equivalence between the mapped controls.

Currently, the CCM is mapped against 35 different standards, best practices, laws and regulations. Some of them are national requirements specific for the cloud. Others are generic information security or privacy regulations, international standards, and sector-specific best practices.

Currently, CCM maps to the following standards:

- ISO/IEC 27001-2013
- ISO/IEC 27002:2013
- ISO/IEC 27017-2015
- ISO/IEC 27018-2014
- AICPA TSC 2017
- AICPA TSC 2014
- C5 BSI Germany
- Canada PIPEDA
- COBIT 5.0
- COBIT 4.1
- COPPA
- CSA Enterprise Architecture
- CSA Guidance V3.0
- ENISA IAF
- PCI DSS v3.0
- PCI DSS v3.2
- FedRAMP Security Controls (Final Release, Jan 2012) - Low Impact Level
- FedRAMP Security Controls (Final Release, Jan 2012) - Moderate Impact Level
- HITRUST CSF v8.1
- 95/46/EC – European Union Data Protection Directive
- FERPA
- GAPP (Aug 2009)
- HIPAA/HITECH Act
- BITS Shared Assessments SIG v6.0
- ITAR
- Jericho Forum
- Mexico, Federal Law on Protection of Personal Data

- NERC CIP
- NIST SP 800-53 R3
- NIST SP 800-53 R4
- NZISM
- ODCA UM: PA R2.0
- CIS-AWS-Foundation v1.1
- Shared Assessments 2017 AUP
- IEC 62443-3-3:2013

Given their global footprint, many CSPs are often required to comply with several of these standards at the same time. For instance, a provider willing to offer services to both the US federal government and the German federal government would need to satisfy the requirements of both FedRAMP and BSI C5. It is therefore useful to have a tool that creates a link between all the various standards, especially considering that frequently the requirements between these various standards overlap.

Based on the results of an analysis conducted by CSA in the context of the European Commission funded project EU-SEC,[203, 204] the controls included in CCM are able to satisfy more than 80% of the requirements included in ISO/IEC 27001, ISO/IEC 270017, ISO/IEC 27018, BSI C5, ENISA Minimum Security Measures for DSPs, AICPA TSC 2016, France's SecNumCloud, and other national standards from Spain and Slovenia. If a company needed to implement each standard from scratch, it would go through considerable duplication of effort, which is an additional cost. Linking the CCM with other international standards and regulations streamlines security and compliance for providers and customers.

Mappings are useful to guide organizations in adhering to new standards or regulations. Take the example of a CSP that would like to expand beyond its APAC and European operations to enter the US market. It might be useful, or even necessary, to show adherence to HIPAA and FedRAMP. In this situation the CCM mapping becomes very useful to help understand how the organization's existing security program and internal control framework can be leveraged to satisfy the new requirements and bridge gaps where they exist.

3.7.2 Gap Analysis and Reverse Mappings

The CCM mapping approach has evolved over time. In the beginning, the mapping constituted a way to link a CCM control to one or several controls of a target standard. (See **figure 3.15**.)

[203] *Op cit* EU-SEC
[204] *Op cit* EU-SEC, "D 1.2 Security and Privacy Requirements and Controls"

Figure 3.15—Evolution of CCM Mapping Approach—Early Stage

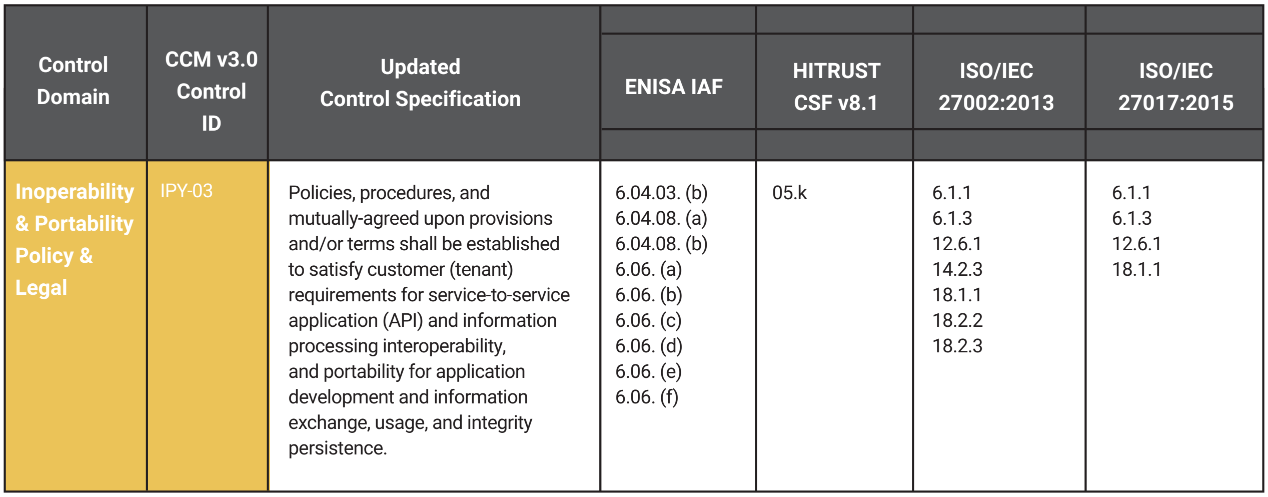

Control Domain	CCM v3.0 Control ID	Updated Control Specification	ENISA IAF	HITRUST CSF v8.1	ISO/IEC 27002:2013	ISO/IEC 27017:2015
Inoperability & Portability Policy & Legal	IPY-03	Policies, procedures, and mutually-agreed upon provisions and/or terms shall be established to satisfy customer (tenant) requirements for service-to-service application (API) and information processing interoperability, and portability for application development and information exchange, usage, and integrity persistence.	6.04.03. (b) 6.04.08. (a) 6.04.08. (b) 6.06. (a) 6.06. (b) 6.06. (c) 6.06. (d) 6.06. (e) 6.06. (f)	05.k	6.1.1 6.1.3 12.6.1 14.2.3 18.1.1 18.2.2 18.2.3	6.1.1 6.1.3 12.6.1 18.1.1

Source: Cloud Security Alliance, Cloud Controls Matrix (CCM) Working Group

Such an approach provided a way to create a link between controls, but it did not specify a few other important details including:

- Which portion of the CCM control was covered by which controls in the target framework in the case of one-to-many mapping?
- Were the multiple requirements included in the CCM control objectives entirely satisfied by the controls in the target framework?
- If gaps existed, where were those gaps?

Another limitation was that the CCM mapping was unidirectional—i.e., it would indicate the connection between a CCM control and one or many controls in the target framework, but it would not provide the target framework perspective:

- Control 123 in framework X corresponds to which CCM controls?
- Are there any gaps?

To overcome these limitations, CSA introduced the concepts of reverse mapping and gap analysis in 2017.

3.7.3 Mapping Methodology

The CCM mapping methodology is composed of three phases:

1. **Preparation**—This includes the initial preparatory actions, such as defining the scope and creating the execution instructions, leadership and work packages.
2. **Execution**—During this phase, the actual work of mapping and gap analysis is conducted. The execution phase includes completion of the controls mapping, gap identification and gap analysis.

3. **Peer review and publication**—This is the closing phase, when the peer review of the final draft takes place according to the CSA Research Lifecycle.[205]

A more detailed explanation about CCM mapping methodology can be found in the CSA document "Methodology for the Mapping of the Cloud Controls Matrix."[206]

Mapping and Reverse Mapping

This section explains CCM mapping and reverse mapping.

Mappings—CCM as Base to Other Target

Each control in the CCM is initially matched to controls in another framework to make an equivalency determination. This approach considers which CCM controls are associated with the control in other established frameworks, and to what degree they are equivalent to each other, thus estimating the extent of new efforts necessary to align to the CCM requirements when another framework is already in use (**figure 3.16**).

Example: An example use case is that of conducting a mapping and gap analysis exercise using CSA CCM as the base framework (CCM controls are positioned on the left side of the mapping tool) and BSI C5 as the target framework (C5 mapped controls are positioned on the right side of the mapping tool). In this mapping configuration, the C5 controls are mapped against each CCM control. The gap analysis shows C5 controls meeting CCM control requirements partially, fully or not at all.

Figure 3.16—CCM Control Mapping Example

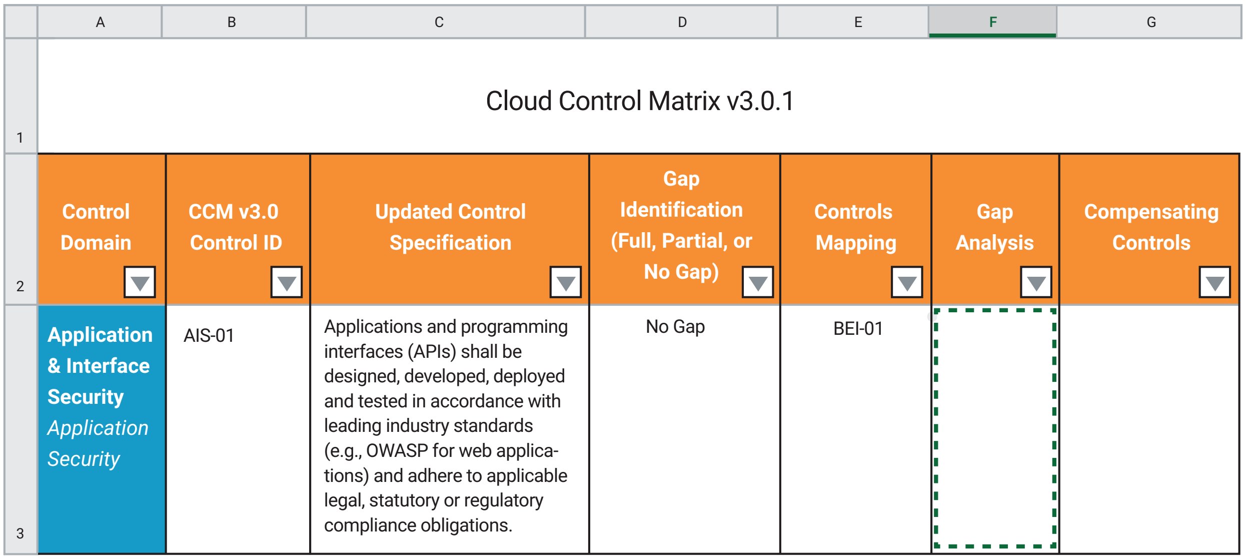

	A	B	C	D	E	F	G
1	Cloud Control Matrix v3.0.1						
2	Control Domain	CCM v3.0 Control ID	Updated Control Specification	Gap Identification (Full, Partial, or No Gap)	Controls Mapping	Gap Analysis	Compensating Controls
3	**Application & Interface Security** *Application Security*	AIS-01	Applications and programming interfaces (APIs) shall be designed, developed, deployed and tested in accordance with leading industry standards (e.g., OWASP for web applications) and adhere to applicable legal, statutory or regulatory compliance obligations.	No Gap	BEI-01		

Source: Cloud Security Alliance, Cloud Controls Matrix (CCM) Working Group

[205] Cloud Security Alliance, "Research Lifecycle," https://cloudsecurityalliance.org/research/lifecycle/
[206] Cloud Security Alliance, "Methodology for the Mapping of the Cloud Controls Matrix (CCM)," 2018, https://downloads.cloudsecurityalliance.org/assets/research/cloud-controls-matrix/ccm-mapping-methodology.pdf

Reverse Mappings—Other Target as Base to CCM

Conversely, a reverse mapping uses another framework as the initial base to find equivalent controls addressed in the CCM. In a reverse mapping exercise, each control in the candidate framework is accounted for in the CCM when possible. The base framework/starting point used in a reverse mapping exercise is the candidate framework. Besides the shift in perspective, all other processes—such as the mapping processes, gap analysis and integration of new requirements—remain the same (**figure 3.17**).

Example: An example use case is that of conducting a reverse mapping and gap analysis exercise using BSI C5 as the base framework (C5 controls are positioned on the left side of the mapping tool) and CSA CCM as the target framework (CCM mapped controls are positioned on the right side of the mapping tool). In this mapping configuration, the CCM controls are mapped against each C5 control. The gap analysis results in having the CCM controls meeting the C5 control requirements partially, fully or not at all.

Figure 3.17—CCM Control Mapping Example: Compliance Controls Catalog (C5)

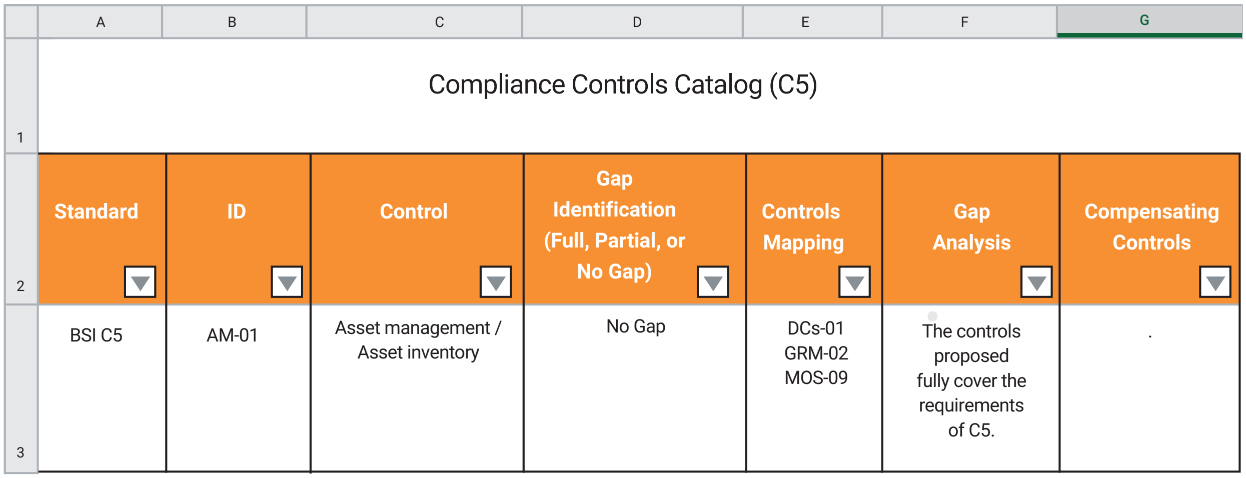

Compliance Controls Catalog (C5)

Standard	ID	Control	Gap Identification (Full, Partial, or No Gap)	Controls Mapping	Gap Analysis	Compensating Controls
BSI C5	AM-01	Asset management / Asset inventory	No Gap	DCs-01 GRM-02 MOS-09	The controls proposed fully cover the requirements of C5.	.

Source: Cloud Security Alliance, Cloud Controls Matrix (CCM) Working Group

Reporting Gap Identification and Analysis

Gap summaries identify full gaps and partial gaps. Full gaps indicate there are no matching controls in the candidate framework. Partial gaps indicate one or more matching controls do exist in the candidate framework, but they do not fully cover the extent of the specific control as defined in the base framework. A complete gap analysis can help inform planning efforts when determining the suitability of extending existing compliance documentation to match another framework. For the purposes of the CCM, a gap analysis specifically lists and explains the gaps between controls in the CCM and another framework.

Mapping Process and Completeness

The goal of the mapping process is to identify a semantic equivalence between CCM controls and those included in other frameworks. Elements that can be used to evaluate equivalence include identifying whether both controls have the following:

- Matching domain names
- Matching security control names

- Matching security control requirements and keywords, in both the other framework and the CCM
  - Each security-related keyword found in a control requirement of the other framework should be investigated to determine if a key word of semantic equivalence exists in the CCM.
- Content matching between security controls identified in both frameworks
  - This approach considers content in a more thorough manner, and successful matches are not as easily identifiable as in the keyword search approach.

Tip: A mapping that is semantically equivalent refers to a scenario in which two or more controls between different frameworks completely cover each other in terms of scope. These controls are said to be semantically equivalent to each other.

Gap Identification and Analysis: No Gap, Partial Gap and Full Gap Definitions

Depending on the scope and objective of the project, a gap identification and analysis might be conducted after the initial mapping for remaining items that were not considered equivalent. A gap identification is essentially an analysis of two or more frameworks that seeks to determine the level of semantic equivalence between frameworks. During the gap identification process, three potential outcomes are possible: no gap, partial gap and full gap. Consider the following to identify which of those three scenarios applies:

- **No Gap**—For a specific CCM control and its requirements, there is an equivalent control or set of controls (in the candidate framework) that fully satisfies the requirements of a corresponding control in the CCM.
- **Partial Gap**—For a specific CCM control and its requirements, there is a control or set of controls (in the candidate framework) that do not fully satisfy the requirements of the corresponding control in the CCM. For the partial gap case to hold, there should be at least one control in the candidate framework that is of semantic equivalence to a requirement in the CCM control. The relevant controls in the candidate framework should then be cited in the work package.
- **Full Gap**—For a specific CCM control and its requirements, there is no control or set of controls (in the candidate framework) that is of semantic equivalence. Essentially, this means that the CCM control is completely unaddressed by any control in the candidate framework.

The preceding examples describe an analysis of the gaps that a candidate framework has compared to the CCM. The same logic can be applied to determine whether the CCM has gaps compared with a candidate framework. See **figures 3.18** and **3.19**.

Figure 3.18—CCM Control Mapping Example: ISO 27002

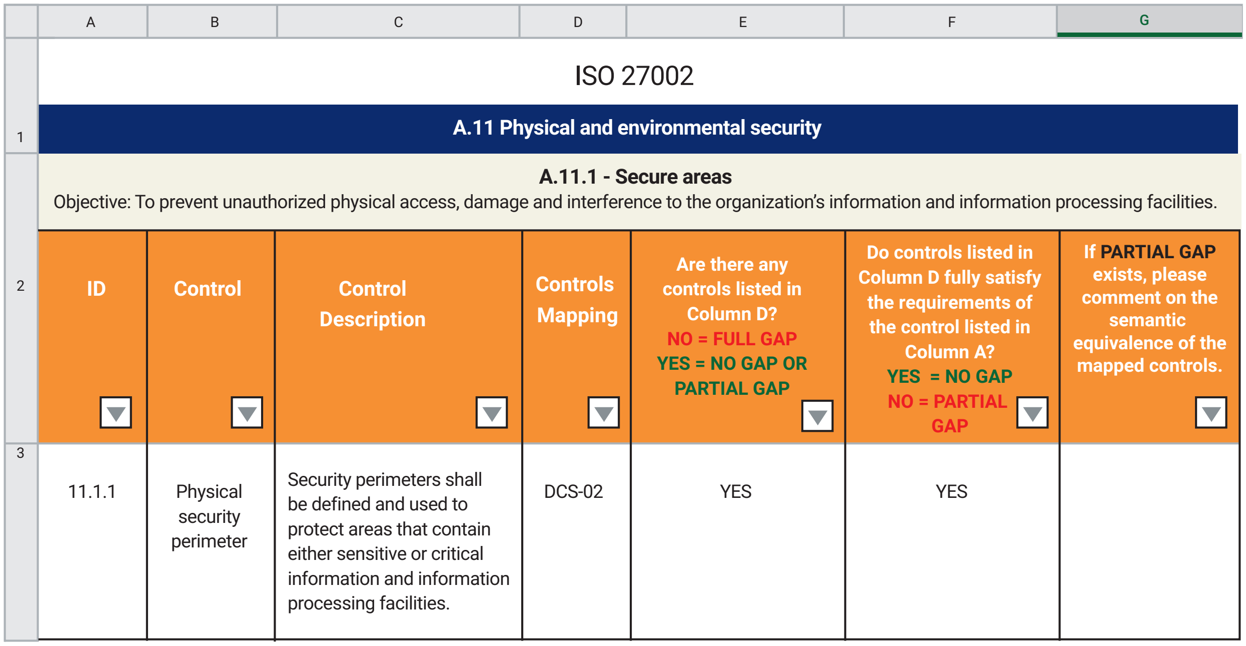

	A	B	C	D	E	F	G
1	ISO 27002 A.11 Physical and environmental security						
2	A.11.1 - Secure areas Objective: To prevent unauthorized physical access, damage and interference to the organization's information and information processing facilities.						
	ID	Control	Control Description	Controls Mapping	Are there any controls listed in Column D? NO = FULL GAP YES = NO GAP OR PARTIAL GAP	Do controls listed in Column D fully satisfy the requirements of the control listed in Column A? YES = NO GAP NO = PARTIAL GAP	If PARTIAL GAP exists, please comment on the semantic equivalence of the mapped controls.
3	11.1.1	Physical security perimeter	Security perimeters shall be defined and used to protect areas that contain either sensitive or critical information and information processing facilities.	DCS-02	YES	YES	

Source: Cloud Security Alliance, Cloud Controls Matrix (CCM) Working Group

Figure 3.19—CCM Key Terms and Definitions

TERMS	DEFINITIONS
Candidate framework	Refers to any recognized set of evaluation criteria (e.g., standard, regulation or code-of-practice), whether international or local technology or industry-specific frameworks.
Domain	A set of related security controls categorized into a specific topic. For example, the CSA Cloud Controls Matrix (CCM) topics relate to the 14 CSA Domains.
Full gap	A similar criterion (control) does not exist in the other framework.
Gap analysis	The breakdown of additional indicators and efforts needed to bridge a control requirement from one framework to another.
Mapping	CCM controls are individually referenced to a non-CCM framework. A mapping can be one-to-one or one-to-many.
No gap	The relevant control requirement in one framework is fully equivalent to another framework.
Partial gap	Controls in two frameworks are similar, but not fully equivalent.
Reverse mapping	The same as mapping, except the starting point is a non-CCM (candidate) framework rather than the CCM.
Semantic equivalence	When security controls between two frameworks are determined to share the same meaning, based on context (how they were written and categorized).
Security controls	Technical or administrative safeguards or countermeasures to modify information security risk to an organization.

3.8 Transition From CCM V3.0.1. to CCM V4

Version four (V4) of the CCM is expected to be released in late 2020. The current version three (V3.0.1) was released in 2014 and a few minor updates were issued since then, including additional mappings with a new framework. In 2018 and 2019, CSA published the ISO addendum and the German BSI C5 addendum, to provide users with a gap analysis and compensating controls between the CCM and ISO/IEC 27001, 27002, 27017 and 27018, and C5.

The CCM, like any other cybersecurity control framework, is subject to obsolescence and needs a periodic major update in order to better reflect the evolution of cloud technology, its adoption strategies, changes in the Guidance, new technologies related to the cloud, and changes in the threat landscape. As CSA went through the mappings with other frameworks, it realized that there were several recently published cloud frameworks that were adding controls specifications that were missing in the Cloud Controls Matrix. There are also several technological gaps with the CCM that CSA wanted to close by including new control references. For example, CSA conducted substantial research work on container, microservices and serverless, which is not included in CCM V3.0.1. Thus, there was a substantial need for a major update from version three to version four.

One of the other reasons for moving to the fourth version was making the CCM controls more auditable. The CCM is the reference framework for the CSA STAR program. Over the past few years, auditors who have audited the implementation of CCM controls for this program have provided feedback. The lessons learned from these auditors are currently being used to increase the clarity of the language and make controls more precise and easier to audit.

3.8.1 Understanding the Changes Between V3 and V4

This section describes the changes from CCM version 3 to version 4.

Domain Structure

For version 4, CSA wanted to avoid major changes in the CCM domains structure. There were multiple reasons for this choice:

- The current version 3.0.1 has a structure that is flexible enough to embrace the technical and organizational changes that have occurred in the cloud market in the past six years and should remain valid in the short-to-medium term (three to five years). In other words, the existing structure can be considered future proof.
- The users of any framework, including CCM, appreciate continuity, and evolutionary approaches are typically preferred to a complete upheaval of the framework structure unless a more substantial change is required.
- CCM is the foundational framework of the CSA STAR program, and there is a need to ensure continuity and comparability from one version to the next. Lack of continuity between versions can shock the user and delay the adoption of the new version of the framework because adapting to the changes requires too much work. That said, there will be changes in the names of some domains.

Domain name changes include:

- Domain EKM (Encryption and Key Management) will be renamed Cryptography, Encryption and Key Management.
- The GRM domain (Governance and Risk Management) will be renamed: Governance, Risk and Compliance (GRC)
- Domain AAC (Audit, Assurance and Compliance) will be renamed AnA (Auditing and Assurance), and the Compliance topic will be moved to the GRC domain.

- Domain Data Security & Information Lifecycle (DSI) will become Data Security & Privacy Lifecycle Management (DSP), and its scope extended from security only to security and privacy.
- The Mobile Security domain will have a more general scope and become Cloud End-point Security.

Merging CCM and CAIQ

Until now, the CCM and CAIQ have been published as two separate research documents. In version 4, however, the CCM and CAIQ will be merged. The CAIQ will be an additional column in the CCM framework in order to help simplify the adoption.

Implementation and Auditing Guidelines

The new version will include implementation and auditing guidance documents to facilitate the use of the CCM.

The implementation guidelines provide suggestions, recommendations and examples of how to implement the CCM controls. Given a certain control specification, the document will help explain what should be done to effectively implement and monitor the control, which specific best practices should be followed, what the specific regulations of reference are, and what the differences are when implementing a control from the SaaS-PaaS-IaaS perspective.

The auditing guidelines provide suggestions, recommendations and examples of how to evaluate whether CCM controls have been correctly implemented. Given certain control specifications, the document will provide suggestions on which source of evidence should be checked, how to verify the maturity of the control, and how to test technical and procedural controls.

Transition Policy.

CSA transition policy foresees that organizations using the previous version of CCM as the foundation for their compliance programs will be entitled to an 18-month grace period to transition to the new version. This means that companies with the CSA STAR certification or attestation will still be able to use it for 18 months following the new release before needing to transition to the new version.

3.9 Chapter 3 Knowledge Check

REVIEW QUESTIONS

1. CCM was created to:

   A. Help only large-scale cloud service providers and auditors
   B. Support cloud customers building a third-party risk management program and a cloud audit plan
   C. Support cloud auditors and consultants in streamlining the auditing process to achieve an ISO 27001 certification
   D. Help CSPs mapping organizational, operational and legal requirements to control objectives.

2. Select **ALL** the applicable correct answers related to infrastructure and virtualization:

A. The IaaS provider is responsible for the security of the resource workloads.
B. The hardening of host and guest OS, hypervisor or infrastructure control plan is an important control to implement.
C. The IVS domain aligns with the recommendation and guidelines included in domains 7 and 8 (Infrastructure Security and Virtualization and Containers) of the CSA Security Guidance.
D. Infrastructure systems and network components should be designed to ensure segmentation and segregation.

3. The CAIQ can be used for the following different purposes (select all that apply):

A. Verify, during the service selection process, that a CSP meets the customer's security specifications
B. Monitor that the customer's security controls objectives are met over time
C. Streamline a CSP approach to customer requests for questionnaires during the cloud service assessment process
D. Guide internal auditors during the assessment of the organization's compliance program

4. The CCM Architectural Relevance column describes:

A. The relationship of the question with the CSA Enterprise Architecture
B. With an x, the applicability of a control to the physical infrastructure, network, compute, storage, applications, data or all these areas
C. The applicability to the network, compute, storage, applications, data or all these areas
D. The applicability to the CSA Enterprise Architecture and physical infrastructure, network, compute, storage, applications, data or all these areas

5. The CCM mapping features are useful to (select all the correct options):

A. Guide organizations in adhering to new standards or regulations
B. Create new controls that an organization can add to its security framework
C. Create links between the CCM controls and over 40 standards and regulations
D. Understand the possible gaps between the existing internal control framework of an organization and the over 40 standards and regulations mapped to CCM

6. Select **ALL** the current statements related to CCM V4 structure (select all the correct statements):

A. The GRM domain (Governance and Risk Management) will be merged with AAC (Audit, Assurance and Compliance) and renamed: Governance, Risk and Compliance (GRC).
B. Domain EKM (Encryption and Key Management) will be renamed Cryptography, Encryption and Key Management.
C. The CAIQ will be merged with the Enterprise Architecture to facilitate the user determining the shared responsibility model.
D. Domain AAC (Audit, Assurance and Compliance) will be renamed AnA (Auditing and Assurance), and the Domain Data Security & Information Lifecycle (DSI) will become Data Security & Privacy Lifecycle Management (DSP).

Answers on page 228

Chapter 3 ANSWER KEY

Correct answers appear in **bold** font.

1. A. CCM was created having CSP, customers and consultants/auditors as a target audience.

   B. Supporting cloud auditors and consultants in streamlining the auditing process to achieve an ISO 27001 certification is one of the reasons why CCM was created.

   C. Wrong answer, since CCM is meant to be used to streamline the organization compliance effort, but not to help specifically to achieve the ISO 27001 Certification

   D. Wrong answer, since CCM helps cloud customers in mapping organizational, operational and legal requirements to control objectives

2. A. Wrong answer, since in IaaS the security of the workloads is a customer responsibility

   B. Correct answer. According to the CSA's best practices, the hardening of host and guest OS, hypervisor or infrastructure control plan is an important control to implement.

   C. Correct answer. The IVS domain aligns with the recommendation and guidelines included in domains 7 and 8 (Infrastructure Security and Virtualization and Containers) of the CSA Security Guidance.

   D. Correct answer. According to the CSA's best practices, infrastructure systems and network components should be designed to ensure segmentation and segregation.

3. **A. Correct answer. Verifying, during the service selection process, that a CSP meets the customer's security specifications is one of the purposes of the CAIQ.**

   B. Correct answer. Monitoring that the customer's security controls objectives are met over time is one of the purposes of the CAIQ.

   C. Correct answer. Streamlining CSP's approach to customer requests for questionnaires for service security assessment is one of the purposes of the CAIQ.

   D. Correct answer. Guiding an internal auditor during the assessment of the organization compliance program is one of the purposes of the CAIQ.

4. A. The CCM Architectural Relevance is used to specify whether a certain control is applicable to the physical infrastructure, network, compute, storage, applications, data or all these areas.

   B. The "x" indicates the applicability of a control to the physical infrastructure, network, compute, storage, applications, data or all these areas.

   C. The CCM Architectural Relevance is used to specify whether a certain control is applicable to the physical infrastructure, network, compute, storage, applications, data or all these areas.

   D. The CCM Architectural Relevance is used to specify whether a certain control is applicable to the physical infrastructure, network, compute, storage, applications, data or all these areas.

5. **A. The mapping feature guides organizations in adhering to new standards or regulations.**

   B. The mapping feature is not a support to create new controls, but to link CCM controls with other standards, best practices and regulations.

   C. The mapping feature creates links between the CCM controls and over 40 standards and regulations.

   D. The mapping feature helps the organization understand the possible gaps between an organization's existing internal control framework and the over 40 standards and regulations mapped to CCM.

6. A. The GRM domain (Governance and Risk Management) will be renamed: Governance, Risk and Compliance (GRC), while Domain AAC (Audit, Assurance and Compliance) will be renamed AnA (Auditing and Assurance), and the Compliance topic will be moved to the GRC domain.

B. Domain EKM (Encryption and Key Management) will be renamed Cryptography, Encryption and Key Management.

C. The merge of CAIQ into CCM won't have a direct effect on the definition of the shared responsibility model.

D. Domain AAC (Audit, Assurance and Compliance) will be renamed AnA (Auditing and Assurance), and the Domain Data Security & Information Lifecycle (DSI) will become Data Security & Privacy Lifecycle Management (DSP).

Page intentionally left blank

Chapter 4:

A Threat Analysis Methodology for Cloud Using CCM

A Threat Analysis Methodology for Cloud Using CCM

A Threat Analysis Methodology for Cloud Using CCM

4.1 Learning Objectives

After completing this chapter, learners will be able to:

1. Describe threat analysis essentials.
2. Use the Top Threat Analysis Methodology to analyze attack details.
3. Document attack impacts based on the Top Threat Analysis Methodology.
4. Apply Threat Analysis Methodology for cloud using CCM.
5. Evaluate a Top Threats method use case.

4.2 Overview

This chapter focuses on post-incident analysis with the aim of educating external auditors and cloud customers (internal auditors, risk managers or other functions) on how to assemble, identify and classify the essential components of an incident. The goal is to enable them to perform a meaningful and comprehensive analysis to accomplish the following:

- Determine the governance and compliance program failures (e.g., weaknesses, gaps in the control design or implementation phases).
- Determine the necessary mitigations (what the organization could have done better, and what can be done to mitigate the issues).
- Understand how to interpret past incidents and breaches as indicators of possible future noncompliance (for the auditor, evaluating past events and the mitigation measures taken and not taken often provides valuable insight for recommending future controls).

4.3 Threat Analysis Essentials

The Top Threat Analysis Methodology created by the CSA Top Threats Working Group (WG) comprises a structured process to identify post incident the who, what, why and how contributing factors, the effects of the cyberincident, and what can be done to prevent a similar event in the future.

Post-incident analysis differs from threat analysis or risk analysis (see chapter 1, section 1.4.10 on cloud risk management) in that the latter are normally anticipatory. The incident has not happened, and the goal is to anticipate the likelihood of the incident—how and why it might occur, the probability it could happen, its potential impacts, and what can be done to prevent it.

4.3.1 Definition and Purpose

The Top Threat Analysis Methodology can be defined as a five-step, in-depth examination of an incident to determine the following:

- **Who, what, why, how and impacts of the incident**
  - **Step 1: Attack details**—Threat actors, threats, vulnerabilities; to determine the who, what, why and how the incident happened.
  - **Step 2: Technical impacts**—Confidentiality, integrity and availability; to determine the actual effects on the system and data.

- **Step 3: Business impacts**—Financial, operational, compliance and reputational impacts; to determine the effects on profit, the ability to deliver goods and services, compliance with applicable laws, regulations and industry standards, and any shift in brand value.

- **What can be done to prevent an incident from happening in the future**
  - **Step 4: Controls**—Preventive, detective, corrective controls; to determine what can be done to prevent the incident from happening in the future, to detect if it is currently happening—and if so, corrective controls that can be used to recover with minimal impact. Included here are implementation guidance, i.e., the considerations and approach; to determine how best to implement controls efficiently and effectively.
  - **Step 5: Metrics**—Key performance indicator attributes, i.e., what, how and when to measure; and control effectiveness measurements, i.e., how best to measure the performance of controls to ensure they are achieving their stated control objectives.

4.3.2 Tools

The Top Threat Analysis Methodology uses a variety of tools to analyze and mitigate an incident. Following are a few examples.

- CSA Tools
  - **CSA Top Threats and Survey**—This report provides organizations with an up-to-date, expert-informed understanding of cloud security concerns, so they can make educated risk management decisions regarding cloud adoption strategies.[207]
  - **CSA Top Threats Deep Dive**—This case study attempts to connect all the dots when it comes to security analysis by using anecdotes cited in the Top Threats for its foundation. Each anecdote is presented in the forms of a reference chart and a detailed narrative. The reference chart format provides a synopsis of the attack, from identifying the actor, threats and vulnerabilities to applying end controls and mitigations.[208]
  - **CSA Cloud Controls Matrix (CCM) v 3.0.1**—The CCM is the only metaframework of cloud-specific security controls, mapped to leading standards, best practices and regulations. It provides organizations with the required structure, detail and clarity to guide decisions regarding information security tailored to cloud computing. CCM is used to define mitigating controls.[209]
- Non-CSA Tools
  - **MITRE ATT&CK®**—This is a globally accessible knowledge base of adversary tactics and techniques based on real-world observations that are updated continuously. The ATT&CK knowledge base is used as a foundation for developing specific threat models and methodologies in the private sector, in government, and in the cybersecurity product and service community.[210]
  - **NIST Vulnerability Database (NVD)**—The US government repository of standards-based vulnerability management data represented using the security content automation protocol (SCAP), which is updated continuously. These data enable automation of vulnerability management, security measurement and compliance. The NVD includes databases of security checklist references, security-related software flaws, misconfigurations, product names and impact metrics.[211]
  - **US National Cyber Awareness System**—This system from the US Department of Homeland Security provides alerts with timely information about current security issues, vulnerabilities and exploits.[212]

[207] *Op cit* Cloud Security Alliance, "Top Threats to Cloud Computing"
[208] *Ibid.*
[209] Cloud Security Alliance, "Cloud Controls Matrix v3.0.1," 3 August 2019, https://cloudsecurityalliance.org/artifacts/cloud-controls-matrix-v3-0-1/
[210] The MITRE Corporation, "MITRE ATT&CK," https://attack.mitre.org/
[211] NIST, "National Vulnerability Database," https://nvd.nist.gov/vuln
[212] US Cybersecurity and Infrastructure Security Agency (CISA), "National Cyber Awareness System – Alerts," www.us-cert.gov/ncas/alerts

4.4 Top Threat Analysis Methodology (Part 1)

Part 1 of The Top Threat Analysis Methodology deals with:

- **Attack Details**—Threat actors, threats and vulnerabilities
- **Technical Impacts**—Confidentiality, integrity and availability

The business impacts are financial, operational, compliance and reputational.

4.4.1 Attack Details

This section is devoted to the attack and its details:

- **Threat actors**—The individuals or groups who lead or carry out the attack
- **Threats**—Events or actions that take advantage of vulnerabilities
- **Vulnerabilities**—Defects in a process, system, application or other asset (including users), that are exploited

Threat Actors

A threat actor is the entity, person, group or organization that accomplishes the threat. Threat actors can be categorized as external or internal, and as malicious or non-malicious.

- **External**—Threats that originate outside the organization:
    - **Malicious**—Intentionally performs some action that damages the confidentiality, integrity, or availability of the system or data.
        - Cybercriminals—Technically proficient professionals using phishing, malware, ransomware, botnets, DDoS attacks and other techniques for profit, to steal data or disrupt a target's operations or service
        - Hacktivists—Those with a social or political agenda who deface websites or launch DDoS or other attacks as a means to effect change
    - **Nonmalicious**—Unintentionally performs, or fails to perform, some action that damages the confidentiality, integrity or availability of the system or data
        - Application user—User who repeatedly submits requests without waiting for web application responses, or who unknowingly takes other steps that negatively affect system performance or data integrity
- **Internal**—Threats that originate within the organization
    - **Malicious**—Intentionally performs some action that damages the confidentiality, integrity or availability of the system or data
        - Disgruntled employees/insiders—Employees, agents or subcontractors who use their access and knowledge to personally benefit or punish the company; may involve stealing data or money, or causing a deliberate service interruption
    - **Nonmalicious**—Unintentionally performs, or fails to perform, some required action that damages the confidentiality, integrity or availability of the system or data
        - Careless employees/insiders—Employees, agents or subcontractors who despite proper instruction continue to access Internet sites recklessly, or otherwise fail to follow corporate security practices
        - Untrained employees/insiders—Employees, agents or subcontractors who have not received proper instruction regarding the handling of phishing, malware and ransomware emails, or other corporate security practices

- Poor performing IT staff/insiders—Technically incompetent employees or third parties who do not perform their duties properly (e.g., misconfigure a database or server, leaving the system or data exposed)

Threat Event

Threats are events or actions carried out by a threat actor that can lead to damage to an organization's operations, assets, employees or reputation through unauthorized access, destruction, disclosure or modification of information, or denial of service.

Examples of threats include the following:

- A cybercriminal group sends a spear phishing email to a careless IT developer (one who previously received security awareness training) with a link to a website containing RAR (self-extracting file) executables. The IT developer opens the file, leading to the exploit of a weakness in the operating system or database that allows a remote attacker to execute arbitrary code on the system with the victim's privileges.
- An IT architect is moving to a new company and, without malice, wants to take copies of work for which he has won an industry award. While at home, he downloads confidential architectural documents—in accordance with his designated access rights and the need-to-know-principle—onto his personal laptop.
- A hacktivist organization discovered a polluting company's unprotected database installed by an untrained consultant, containing the locations of unreported contaminated water supply, which could be accessed without any authentication or access control via an open 27017 port.
- A recently fired developer sent a glowing farewell email that was opened by the CISO. The email carried a .wsf file containing detection evasion scripts combining multiple programming languages, allowing it to pass through emulation engines that rely on a single language.

A state-sponsored group compromised Internet of Things (IoT) devices that used default credentials. The attackers used Mirai malware to infect IoT devices to create a botnet, which they then used to launch a DDoS attack on a domain name system (DNS) service provider.

Vulnerabilities

A vulnerability is a deficiency—in a process, system, application, IT asset, system security procedure or internal control—that a threat actor can use to accomplish a threat. Vulnerabilities make the threat actor's goals achievable. They generally relate to missing, weak or misapplied security controls, but they also can result from the following:

- Governance structures with missing or poor risk management
- External relationships with poor processes that have not been risk managed
- Enterprise/information security architectures that are poorly designed

Cloud vulnerabilities can be divided into the following cloud classes.[213]

- **Misconfigurations**
  - Default passwords not disabled
  - Insecure (weak) password policies implemented
  - Files and directories not access-restricted
  - Software end of life (outdated/old versions)

[213] National Security Agency, "Mitigating Cloud Vulnerabilities," 22 January 2020, https://media.defense.gov/2020/Jan/22/2002237484/-1/-1/0/CSI-MITIGATING-CLOUD-VULNERABILITIES_20200121.PDF

- Unpatched software (unmitigated bugs)
- Improper input/output data validation

- **Access**
  - Multifactor authorization (MFA) with strong factors not used
  - Regular reauthentication not required
  - MFA password resets not used
  - Protocols using weak authentication enabled
  - Access between cloud resources not managed using zero trust/SDP
  - Managed authentication between virtual machines not used
  - API keys stored in version control systems
- **Shared tenancy**
  - Misunderstanding the intended use of the underlying isolation technology
  - Poor or no encryption of data at rest and in transit with weak encryption methods
  - Improperly configured, managed and monitored key management systems
  - Dedicated, bare-metal instances not used for sensitive workloads
- **Supply chain**
  - Poor vendor (hardware, software, consulting, service) acquisition processes
  - Insufficient SLA requirements and protections
  - Lack of a process to monitor and enforce vendor relationships (e.g., to perform an audit)
  - Inadequate encryption as data passes through the supply chain

4.4.2 Technical Impacts

The Top Threat Analysis Methodology considers the technical impacts of an attack on a system and its data that compromise confidentiality, integrity and availability (**figure 4.1**): unauthorized disclosure of information, altered information, and untimely and unreliable access to information.

Figure 4.1—CIA Triad

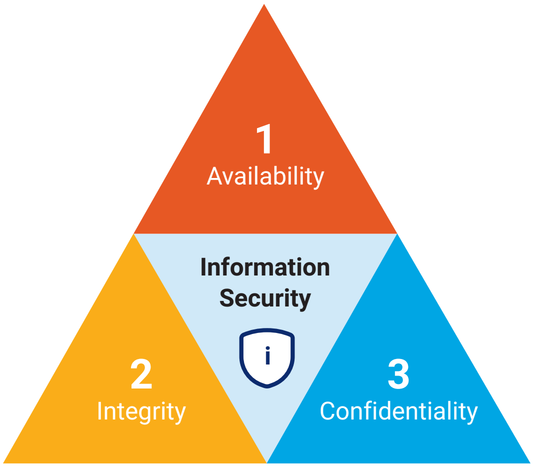

Confidentiality

Confidentiality ensures that information is protected from unauthorized access or disclosure.

Example: An insider who has accepted a job at a competitor downloads new product information onto a laptop before taking the laptop home and downloading the information onto a home computer. The disclosure of technical information to a competitor for a new product might be considered severe (see the confidentiality impact table in **figure 4.2**) if it significantly reduced the company competitive advantage.

An organization impact table is based on likelihood and impact determined when carrying out its risk management procedures (see section 1.4.10 on cloud risk management). **Figure 4.2** shows a confidentiality impact table.

Figure 4.2—Confidentiality Impact Table

Level	Definition
Low	Unauthorized disclosure of information results in limited damage to the organization operations, assets and human resources.
Medium	Unauthorized disclosure of information results in serious damage to the organization operations, assets and human resources.
High	Unauthorized disclosure of information results in severe or catastrophic damage to the organization operations, assets and human resources.

Integrity

Integrity ensures information is complete, accurate and up-to-date and is not subject to unauthorized modification or destruction, and it ensures its non-repudiation and authenticity.

This applies to data in storage, during processing and in transit.

Example: An inexperienced IT consultant new to a boutique investment bank misconfigures one of the logical storage units granting unrestricted access and change rights. This storage unit contains proprietary financial formulas used by in-house bankers, corporate investment departments and high-net-worth individuals who perform their own analyses. A hacker discovers the unprotected storage unit and decides to delete the storage, causing a halt to the investment bank trading operations. This incident might be considered severe (see the integrity impact table) depending on the investment bank losses and those of its customers.

The organization impact table is based on likelihood and impact determined when carrying out its risk management procedures (see section 1.4.10). **Figure 4.3** shows an integrity impact table.

Figure 4.3—Integrity Impact Table

Level	Definition
Low	Unauthorized modification or destruction of information results in limited damage to the organization operations, assets and human resources.
Medium	Unauthorized modification or destruction of information results in serious damage to the organization operations, assets and human resources.
High	Unauthorized modification or destruction of information results in severe or catastrophic damage to the organization operations, assets and human resources.

Availability

Availability refers to ensuring that information, systems, facilities, networks and computers are available to authorized individuals or groups when they need to access them.

Example: A phishing campaign leading to a ransomware infection that prevents a bank from accessing 10,000 high-value customer accounts unless a ransom is paid might be considered severe (see the availability impact table in **figure 4.4**) depending on the length of unavailability and customers' losses, and the bank loss of customers.

The organization impact table is based on likelihood and impact determined when carrying out its risk management procedures (see section 1.4.10). **Figure 4.4** shows an availability impact table.

Figure 4.4—Availability Impact Table

Level	Definition
Low	Unavailability of a small set of information or an information system resulting in limited damage to the organization operations, assets or human resources.
Medium	Restricted access to or use of information or an information system results in serious damage to the organization operations, assets or human resources.
High	Restricted access to or use of information or an information system results in severe or catastrophic damage to the organization operations, assets or human resources.

4.4.3 Business Impacts

The ultimate impact of any incident that can affect the organization can be classified into four categories: financial, operational, compliance and reputational.

Financial Impacts

Financial impacts concern the direct increase in costs that result from an incident, including:

- **Compensation**—Revenue from partners and customers may be lost or additional costs incurred as a result of the organization's security incident.
- **Technical investigations**—Security incidents may generate significant investigatory costs due to the complexity of the environment where the investigation takes place and the need for expert resources to conduct the investigation.
- **Ransomware**—Amounts may have to be paid to an attacker to allow an enterprise to regain access to its system and or data.
- **System upgrades**—Upgrades may be required for the affected system, software or services, or new systems, software or services may have to be purchased to improve the enterprise's overall security posture and reputation in the marketplace.
- **Insurance premium increases**—Security incidents may result in significant costs that may be covered under a cyber risk or other insurance policy. Insurance premiums can be expected to rise as a result, and it may be necessary to hire law firms or other third parties to manage the claims process.
- **Litigation costs**—It may be necessary to recover costs from perpetrators or to defend against supplier or customer lawsuits.
- **Increased financing costs**—A perceived or real incident can increase the interest rate the company pays for its financing, or result in creditors seeking to modify or terminate financing agreements.

- **Reduced investments**—Additional costs incurred due to damages can reduce the expectations of existing and potential shareholders and investors, thus lowering the stock price and making raising debt or equity investment more difficult.
- **Severance pay and other staff termination costs**—An incident may require the affected company to terminate staff, which adds costs through severance packages and related benefits.
- **Recruitment costs**—Recruiting new, more capable staff can require significant hire costs.

Operational Impacts

Operational impacts concern disruptions to business processes, systems and data. These impacts also result in financial impacts. Following are examples of such operational impacts.

- **Reduced product/service sales**—Security incidents can bring sales and marketing systems to a halt, resulting in a delay or inability to process customer orders or manage customer relationships.
- **Production/service delays**—Security incidents can bring public-facing production and service systems down. Production or service orders may not be processed.
- **BOMs (bills of materials) may be corrupted**—A security incident may not allow production according to specifications.
- **Production/service planning files may be corrupted**—Systems affected by a security incident may be unable to process the right products or provide services at the right time.
- **Product/service quality**—Quality systems can be brought down by security incidents, preventing necessary checks from taking place and negatively affecting customer orders or relationships.
- **Product/service delivery**—Logistics systems can be slowed or brought to a halt by security incidents. Additionally, products and services might be delivered to the wrong location at the wrong time.
- **New product/service introduction delays**—Security incidents can delay the launch of new products or services by interfering with the organization's ability to bring production or service supply chains online in time to satisfy the customer.
- **Production/service reporting systems**—A compromise caused by a security incident may not allow correct manufacturing, financial, human resource or other reporting, hampering decision making.

Compliance Impacts

Compliance impacts result from not acting in accordance with applicable laws and regulations. Following are examples of such impacts.

EU GDPR

The General Data Protection Regulation (GDPR)[214] is a legal framework that sets requirements for the collection and processing of the personal data of European Union citizens.

Lack of compliance with the GDPR might impact an organization in several ways:

- Regulatory investigations and fines
- Litigation against affected individuals
- Litigation against other third parties
- Need to hire attorneys and experts to defend against regulatory investigations and litigation
- Costs to improve the legal and compliance function

[214] European Commission, "Data protection in the EU," https://ec.europa.eu/info/law/law-topic/data-protection/data-protection-eu_en

The GDPR imposes fines of up to 10 million euro, or 2% of a firm's worldwide annual revenue from the preceding financial year (whichever amount is higher), for violating the articles governing controllers and processors, certification bodies, and monitoring bodies.

Fines of up to 20 million euro, or 4% of the firm's worldwide annual revenue from the preceding financial year (whichever amount is higher), apply for violations of the articles governing the basic principles for processing, the conditions for consent, the data subjects' rights, and transfer of data to an international organization or a recipient in a third country.

HIPAA

The Health Insurance Portability and Accountability Act (HIPAA) of 1996[215] requires HIPAA-covered entities to safeguard the protected health information (PHI) of patients, and to strictly control when PHI can be divulged, and to whom.

Criminal HIPAA violations include theft of patient information for financial gain and wrongful disclosures with intent to cause harm.

The tiers of criminal penalties for HIPAA violations:

- **Tier 1**—Reasonable cause or no knowledge of violation—up to one year in jail
- **Tier 2**—Obtaining PHI under false pretenses—up to five years in jail
- **Tier 3**—Obtaining PHI for personal gain or with malicious intent—up to 10 years in jail

Reputational Impacts

Reputational impacts are related to perceptions of the company as a whole—but leaning toward internal factors, including management issues and brand value—and stakeholders' perceptions of products, services and processes owned, licensed or provided by an organization. In practice, brand value and reputation are often used interchangeably. These impacts result in financial consequences. Following are examples of such reputational impacts.

- **Damaged public perception**—Public perception can be negatively affected if an enterprise's management or employees are seen as negligent for allowing the incident to occur. This perception can be compounded if public handling and resolution of the incident and its effects are viewed as incompetent.
- **Damaged customer relationships**—Particularly in heavily regulated industries with sensitive data, such as healthcare and finance, security of information and systems is crucial to customer comfort.
- **Damaged supplier relationships**—Security incidents may cause a supplier whose system is connected to the company to be concerned about damage to its own security and reputation from being associated with the breached company.
- **Reduced business opportunities**—Potential customers and suppliers or other partners may balk at aligning with the compromised operations of another company.
- **Recruitment difficulties**—Security incidents, especially highly publicized ones with an inadequate or inappropriate organizational response, can make it difficult to attract and retain top employees who may not want their personal brand tarnished.
- **Key staff loss**—Employees, especially key employees, who have resigned, been made redundant, or been redistributed within the organization due to a security incident might have knowledge of the company's communication channels and access to contacts. How employees communicate with customers is significant. If

[215] HIPAA Journal, "What are the Penalties for HIPAA Violations?" 24 June 2015, www.hipaajournal.com/what-are-the-penalties-for-hipaa-violations-7096/

the number of resignations or redistributions is substantial, it sends a message to the outside world that all is not right with the organization.

- **Media scrutiny**—Members of the media know that fear and disasters draw attention. If there is even a hint of an issue, media outlets will rush to cover the story. If it is a large company with many small customers, then it is David vs. Goliath, and the scrutiny intensifies.

4.5 Top Threat Analysis Methodology (Part 2)

The second part of the Top Threat Analysis Methodology deals with what happens after the attack details are known. Threat actors, threats and vulnerabilities have been identified, the technical impacts (confidentiality, integrity and availability) have been determined, and the business impacts (financial, operational, compliance and reputational) have been assessed. What comes next? Mitigation.

4.5.1 Mitigating Controls

According to ISACA,[216] The IIA[217] and ENISA[218] controls exist to prevent, detect or correct the effects of incidents. Their common objectives are to ensure that such incidents never occur again, or that the chances of their occurring are reduced and their potential impact minimized, or that if a similar incident should occur, it will be discovered and addressed in a timely way to minimize or at least reduce the damage. If the incident is not prevented, and the system and data are damaged, appropriate controls can help ensure that the system and data are recoverable so that business may continue.

Controls are commonly classified as follows (**figure 4.5**):

- **Technical**—Controls implemented and executed by the system
- **Administrative**—Control policies and procedures to guide human behavior
- **Physical**—Systems or people that limit or monitor access and ensure availability of the system

Figure 4.5—Summary of Preventive, Detective and Corrective Controls

	Preventive	Detective	Corrective
Technical controls	• Intrusion prevention system • Antivirus • Firewalls • Patches • Configuring IAM policies	• Intrusion detection system • Vulnerability scans • Configuring events rules	• Backup and restore • Reboot systems • Configure incident rules
Administrative controls	• Security awareness training • Segregation of duties • BYOD policies/training • Patch management	• Reviewing logs • Auditing	• Business continuity • Incident response
Physical controls	• Guards • Fences • Locks	• Door alarms • Fire alarms • Motion setection • CCTV	• Renew access cards • Fire extinguisher

[216] ISACA, "Glossary," https://www.isaca.org/resources/glossary
[217] The Institute of Internal Auditors (IIA), "Global Technology Audit Guide (GTAG) 1: Information Technology Risk and Controls, 2nd Edition," https://na.theiia.org/standards-guidance/recommended-guidance/practice-guides/Pages/GTAG1.aspx
[218] Dekker, M.; C. Karsberg; "Technical Guidelines on Security Measures," ENISA, November 2013

Preventive Controls

In Top Threat Analysis Methodology terms, preventive controls exist to stop a threat from coming into contact with the system or asset.

- Technical controls, such as MFA (multi-factor authorization), deny connection to the system or data pending authentication.
- Administrative controls, such as BYOD (bring your own device) user policy training, helps ensure that devices compatible with the system are securely attached.
- Physical controls, such as guards and badges, prevent unauthorized access to the grounds of a data center.

Detective Controls

Detective controls identify an incident in progress or uncover one that has already achieved its objective.

- Technical controls, such as IDSs (intrusion detection systems), monitor a network for malicious activity or policy violations.
- Any malicious activity or violation is typically reported or collected centrally using a SIEM (security information and event management) system.
- Administrative controls, such as reviewing logs, can help uncover suspicious access or activity leading to discovery of an incident.
- Physical controls, such as motion detection systems and closed-circuit television cameras, can detect the presence of an intruder after entry has occurred.

Corrective Controls

Corrective controls exist to restore the system or process back to its state prior to the incident.

- Technical controls, such as backup restore, ensure that a system can be restored to its normal state prior to the incident or as close to it as possible.
- Administrative controls, such as incident-response plans, ensure that personnel know how to coordinate a timely and proper response to incidents that require restoration of systems or data.
- Physical controls, such as renewing access cards (which includes canceling the old cards), restores users to their positions prior to an incident.

Implementation Guidance

This section provides guidance concerning the implementation of controls selected to prevent, detect or correct the effects of an incident, as a result of a Top Threat Methodology Deep Dive analysis. There are several ways to implement any control. However, configuration, installation and operation must be considered to ensure the control is properly designed and implemented for its application situation. For illustration, consider the following examples using CCM Controls EKM-04 and IAM-07:

EKM-04 considerations include (**figure 4.6**):

- What strength of control to implement depending on use/risk
- Whether to use an in-house or outsourced control
- What limitations to place on the application of the controls
- Whether to single-purpose the control

Figure 4.6—Encryption and Key Management Domain

Storage and Access	**EKM-04**	Platform and data-appropriate encryption (e.g. AES-256) in open/validated formats and standard algorithms shall be required. Keys shall not be stored in the cloud (i.e., at the cloud provider in question) but maintained by the cloud consumer or trusted key management provider. Key management and key usage shall be separated duties.

Implementation

Ensure that the following is implemented:

- Encrypt using sufficiently durable encryption strengths, such as AES-256. Use open, validated formats and avoid proprietary encryption formats wherever possible. Proprietary encryption algorithms are unproven and easily broken.
- Maintain their own keys or use a trusted cryptographic service from a source that currently maintains such a service.
- When encrypting data in database, do not encrypt primary keys or indexed columns.
- In general, a single key should be used for only one purpose (e.g., encryption, authentication, key wrapping, random number generation, or digital signatures).

Source: Cloud Security Alliance, Cloud Controls Matrix (CCM) Working Group

IAM-07 considerations include (**figure 4.7**):

- Where to implement the control and when to use it
- Alternative parties to ensure resilience
- Sharing of access to information and services
- Who can declare a condition prior to initiating a control
- Who can initiate controls
- Who can share control knowledge

Figure 4.7—Identity and Access Management Domain

Third-Party Access	**IAM-07**	The identification, assessment, and prioritization of risk posed by business processes requiring third-party access to the organization information systems and data shall be followed by coordinated application of resources to minimize, monitor, and measure likelihood and impact of unauthorized or inappropriate access. Compensating controls derived from the risk analysis shall be implemented prior to provisioning access.

Implementation

Ensure that the following is implemented:

- Monitor service continuity with upstream provider in the event of provider failure
- Have more than one provider for each service the organization depends on
- Provide access to operational redundancy, continuity summaries and services
- Provide a tenant-triggered failover option
- Share business continuity and redundancy plans with tenants

Source: Cloud Security Alliance, Cloud Controls Matrix (CCM) Working Group

4.5.2 Metrics

Performing a Top Threat Analysis results in identifying controls designed to detect or correct an incident to prevent it from occurring in the future, and then implementing the agreed-upon controls. After controls are identified and put in place, it is necessary to develop KPIs (key performance indicators) that are generated as a product of control performance. Control effectiveness measurement is the monitoring and testing of the controls.

Key performance indicators must enable measurements with the following minimum attributes:

- **Relevant**—Support business and security goals
- **Quantifiable**—Result in numbers, percentages, averages
- **Understandable**—Clear to all stakeholders
- **Evidenced**—Based on supporting data
- **Timely**—Gathered, summarized and interpreted while useful
- **Actionable**—Potential for corrective steps
- **Accurate**—Necessary for optimal actions
- **Ownership**—Responsibility prerequisite for accountability

Control effectiveness measurements concern monitoring and testing with the following attributes:

- **Continuous**—Few or no delays
- **Alerts**—Generate notices to the right stakeholders
- **Multiple notifications**—Provided frequently via mail, phone, in-person
- **Graphic visualization**—Facilitates data analysis
- **Comprehensive**—Relevant access to the underlying data necessary for investigation
- **Forensic**—Enables a detailed technical forensic review if necessary
- **Resolutive**—identifies an action leading to a resolution

Figure 4.8 shows an example of metric/monitoring.

Figure 4.8—Metric/Monitoring Example

Name	Remote Access Control Measure
Measure	Percent of remote access points used to gain unauthorized access
Formula	Number of remote access points used to gain unauthorized access/access points * 100
Target	.05%
Collect/Report frequency	Continuously
Responsible	• Information Owner: Computer Security Incident Response Team (CSIRT) • Information Collector: System Admin or Information System Security Officer (ISSO) • Information Customer: Chief Information Officer (CIO)
Data source	Source Incident database, audit logs, network diagrams, IDS logs and alerts
Report format	Stacked bar chart, with drill down to access points and source data

Source: National Institute of Standards and Technology (NIST), *Performance Measurement Guide for Information Security*, July 2008, https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-55r1.pdf

4.6 Top Threats Analysis Method Use Case

This section presents a use case demonstrating the Top Threats Analysis Method and represented in **figure 4.9**.

Figure 4.9—Top Threats Analysis Method Use Case: Capital One

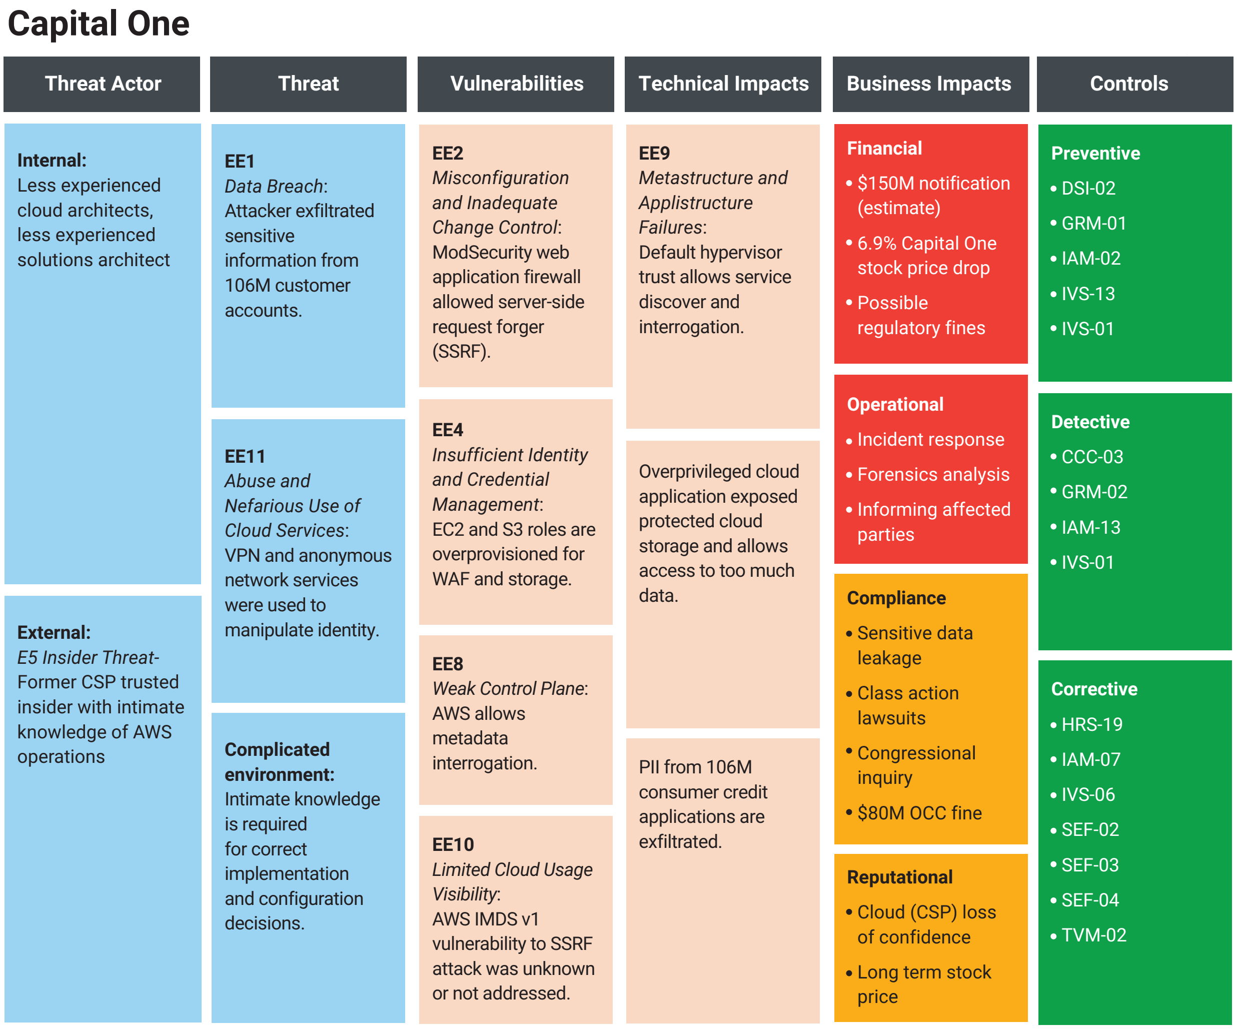

Capital One

Threat Actor

Internal:
Less experienced cloud architects, less experienced solutions architect

External:
E5 Insider Threat-
Former CSP trusted insider with intimate knowledge of AWS operations

Threat

EE1
Data Breach:
Attacker exfiltrated sensitive information from 106M customer accounts.

EE11
Abuse and Nefarious Use of Cloud Services:
VPN and anonymous network services were used to manipulate identity.

Complicated environment:
Intimate knowledge is required for correct implementation and configuration decisions.

Vulnerabilities

EE2
Misconfiguration and Inadequate Change Control:
ModSecurity web application firewall allowed server-side request forger (SSRF).

EE4
Insufficient Identity and Credential Management:
EC2 and S3 roles are overprovisioned for WAF and storage.

EE8
Weak Control Plane:
AWS allows metadata interrogation.

EE10
Limited Cloud Usage Visibility:
AWS IMDS v1 vulnerability to SSRF attack was unknown or not addressed.

Technical Impacts

EE9
Metastructure and Applistructure Failures:
Default hypervisor trust allows service discover and interrogation.

Overprivileged cloud application exposed protected cloud storage and allows access to too much data.

PII from 106M consumer credit applications are exfiltrated.

Business Impacts

Financial
- $150M notification (estimate)
- 6.9% Capital One stock price drop
- Possible regulatory fines

Operational
- Incident response
- Forensics analysis
- Informing affected parties

Compliance
- Sensitive data leakage
- Class action lawsuits
- Congressional inquiry
- $80M OCC fine

Reputational
- Cloud (CSP) loss of confidence
- Long term stock price

Controls

Preventive
- DSI-02
- GRM-01
- IAM-02
- IVS-13
- IVS-01

Detective
- CCC-03
- GRM-02
- IAM-13
- IVS-01

Corrective
- HRS-19
- IAM-07
- IVS-06
- SEF-02
- SEF-03
- SEF-04
- TVM-02

Source: Cloud Security Alliance, Top Threats Working Group

4.6.1 Attack Details (Columns 1, 2, 3)

This section is devoted to the attack and its details. These are: threat actors, the threats and the vulnerabilities.

The attack detail:

- **Actor**—Former engineer of AWS with insider knowledge on platform vulnerabilities gained credentials from a misconfigured web application to extract sensitive information from protected cloud folders.
- **Attack**—Open-source anonymity network (Tor) and VPN service (iPredator) hide attacker. Misconfigured ModSecurity WAF used by Capital One with its AWS cloud operations relayed AWS cloud metadata services including credentials to cloud instances. Overprivileged access given to the WAF allowed the attacker to gain access to protected cloud storage (AWS S3 buckets) with the ability to read data sync and exfiltrate sensitive information.

- **Vulnerabilities**—A server side request forgery (SSRF) vulnerability on the platform tricked a server (i.e., Capital One WAF) into granting requests from an attacker to access cloud server configurations (i.e., EC2 metadata service) providing credentials for whatever the server had access to.

4.6.2 Technical and Business Impacts (Columns 4 and 5)

This section considers the technical impacts of an attack on a system and its data, i.e., confidentiality, integrity and availability; and business impacts related to finance, operations, compliance and reputation.

- **Technical impacts**
  - **Data breach**—A web application was compromised for IAM credentials to access multiple cloud folders. The cloud folders accessed had read rights to 106 million records of customer information that were exfiltrated.
  - **Data loss**—The data extracted were credit card applications and credit card customer status reports from 2005-2019. Personal Identified Information (PII) from the applications included applicant names, addresses, ZIP codes/postal codes, phone numbers, email addresses, dates of birth and self-reported income. The credit card customer PII and financial records extracted included credit scores, credit limits, balances, payment history, contact information, Social Security Numbers and linked bank accounts. Approximately 140,000 Social Security Numbers and 80,000 linked bank account numbers of secured credit card customers were exfiltrated.
- **Business impacts**
  - **Financial**—Exposure of customer bank account information can lead to loss of customer financials and insurance costs for the banking institution. The impact on 106 million customers lead to an OCC settlement of $80 million for customer credit monitoring, identity restoration services, fraud, or other misuse of customer information. Additional regulatory violations may lead to additional fines. The increase in penalties paid and loss of revenue will impact stock prices.
  - **Operational**—Incident response and additional legal investigation, replacement and retraining of security staff, risk and vulnerability assessments and reconfigurations of applications, and notifications to customers and repairing of damage disrupted normal business operations.
  - **Compliance**—Loss of customer PII in violation of GDPR and other privacy regulations can lead to monetary penalties. Higher-regulated industries, such as financial services, are subject to strict monitoring for customer protection with heavy penalties. Equifax faced $575 million in fines from the US Federal Trade Commission in a 2017 data breach that impacted 147 million customers.
  - **Reputational**—Loss of customer and applicant information is expected to impact Capital One customer and public confidence with revenue decreases over three years following the incident due to fewer customer acquisitions. The breach also resulted in reputational losses internally, with the CISO being reassigned and almost a dozen security professionals at the organization quitting.

4.6.3 Mitigating CCM Controls (Column 6)

This section considers existing controls to prevent, detect or correct the effects of incidents to ensure that the incident does not occur again, or that the chance of it occurring is reduced, or that if the incident occurs it will be discovered in a timely way and necessary steps will be taken to place the entity in the position it was in prior to the incident.

- **Preventive mitigation**
  - **DSI-02–Data Inventory Flows**—Inventory, documentation and maintenance of data flows identify and establish the secure archiving, destruction and disposal of aging customer data.
  - **GRM-01–Baseline Requirements**—Established security requirements prevent deviations from baseline configurations and identify vulnerabilities before implementation and use of an application.

  - **IAM-02–Credential Lifecycle/Provision Management**—Appropriate policies, procedures, processes and measures prevent the overprovisioning of access to excessive cloud folders and sensitive information.
  - **IVS-13–Network Architecture**—Architecture diagrams and data flows are applied for timely detection and response to network penetration and the exfiltration of data.
  - **SEF-01–Contact Authority Maintenance**—Points of contact for regulators and law enforcement are maintained for immediate compliance and preparation for forensic investigation when a breach occurs.
- **Detective mitigation**
  - **CCC-03–Quality Testing**—Quality change control and testing are established for application misconfigurations affecting the confidentiality, integrity and availability of the systems and services.
  - **GRM-02–Data-Focused Risk Assessments**—Data-focused assessments identify the proper and improper use, storage, destruction and access of sensitive data.
  - **IAM-13–Utility Programs Access**—Utility programs identify the AWS server vulnerability to an SSRF attack and restrict the IMDS metadata exploit. (AWS IMDSv2 prevents the occurrence of this type of SSRF attack.)
  - **IVS-01–Audit Logging/Intrusion Detection**—Proper log management for suspicious network behaviors and file integrity anomalies are recorded for investigation in the event of a security breach.
- **Corrective mitigation**
  - **HRS-09–Training/Awareness**—Cloud architecture and data lifecycle management identify misconfigurations, over-permissioned applications, and improper data management processes. Continuous training on cloud platforms and security techniques prepares staff for up-to-date platform features and the latest types of attacks.
  - **IAM-07–Third Party Access**—Assessment of risks posed by third-party access to cloud services identifies over-permissioned and other inappropriate access by a WAF or other applications.
  - **IVS-06–Network Security**—Implements the latest design and configuration techniques to monitor access and behavior from trusted and untrusted connections.
  - **SEF-02–Incident Management, SEF-03–Incident Reporting, SEF-04–Legal Preparation**—Response to an incident, breach notification, and forensics procedures is conducted in a timely manner with impacted customers, third parties, regulatory bodies, and other legally required entities.
  - **TVM-02–Vulnerability/Patch Management**—Identifies vulnerabilities such as SSRF in the IMDS platform, and patch or push for a CSP patch.

4.6.4 Metrics

This section considers key performance indicators (KPIs) that are generated as a product of control performance and control effectiveness measurements, which is the monitoring and testing of the controls.

The metrics include:

- **Key performance indicators**—Misconfiguration scans, cloud architecture expertise, data inventory model, credential provisioning
- **Control effectiveness measurements**—Implementation architecture and data flow diagrams, data storage and disposal archiving, access-control alerting

4.6.5 Key

The following are key takeaways:

- Be aware of cloud service metadata that can be exposed with misconfigurations.
- Over-privileged cloud apps allow access to too much data when compromised.
- Data inventory/lifecycle practices for archiving, disposal and destruction limit data exposure.

Page intentionally left blank

4.7 Chapter 4 Knowledge Check

REVIEW QUESTIONS

1. The Top Threats Analysis Methodology looks at post-incident analysis to educate the external auditors and the cloud customers (internal auditors, risk managers or other functions) on how to assemble, identify and classify the essential component parts of an incident to be able to perform a meaningful and comprehensive analysis that can be used to (select all that apply):

   1. Determine the governance and compliance program failures

   2. Determine the necessary mitigations

   3. Understand how to perform a statistical sampling to verify control effectiveness

   4. Understand how to interpret past incidents and breaches as telltales for possible future noncompliance

   A. 2, 4
   B. 2, 3, 4
   C. 1, 2, and 4
   D. 1, 3

2. The Top Threats Analysis Methodology considers the technical impacts of an attack on a system and its data. Each paragraph below describes a different technical impact. Please choose the answer below which correctly labels each technical impact.

   Technical impact #1: An inexperienced IT consultant new to the boutique investment bank misconfigures one of the local servers granting unrestricted access and change rights. A hacker discovers the unprotected server, notes highly active customers and inserts his bank account number in place of the customers account numbers.
   Technical impact #2: A phishing campaign that allows installation of ransomware prevents access to 10,000 high value customer accounts of a bank unless a ransom is paid might be considered severe depending on the length of unavailability and customers losses as well as loss of customers.
   Technical impact #3: An insider who has accepted a job at a competitor senses he is about to be fired from his current company and downloads new product information onto his laptop before taking the laptop home and downloading the information onto his home computer.

   A. Integrity, Confidentiality, Availability
   B. Confidentiality, Availability, Integrity
   C. Integrity, Availability, Confidentiality
   D. Availability, Integrity, Confidentiality

3. Match the monitoring terms in column A with the monitoring descriptions in column B.

Column A: Monitoring Terms	Column B: Monitoring Descriptions
1.__ Continuous	A. Investigatory friendly
1.__ Graphic visualization	B. Necessary relevant statistics
1.__ Comprehensive	C. Facilitates data presentation and analysis
1.__ Forensic	D. With little or no delays

Answers on page 250

Chapter 4 ANSWER KEY

Correct answers appear in **bold** font.

1. A. Although 2 and 4 are correct statements, this answer is missing statement 1, which is another possible use of post-incident analysis.

   B. This answer is incorrect because: 1) Determine the governance and compliance program failures, and 3) Understand how to perform a statistical sampling to verify control effectiveness, are aspects of audit and not a use of the report. The Top Threats Analysis Methodology lists 1, 2 and 4 as its primary uses.

   C. This answer is correct because it addresses the three primary uses of the Top Threats Analysis Methodology listed in chapter 4, A Threat Analysis Methodology for Cloud using CCM, section 4.1.0, Introduction.

   D. Although 1 is correct, this answer is missing statement 2 and 4, which are other possible uses of post-incident analysis.

2. A. This answer is incorrect because the correct labeling, according to the Top Threats Analysis Methodology, is: Integrity, Availability, Confidentiality.

   B. This answer is incorrect because the correct labeling, according to the Top Threats Analysis Methodology, is: Integrity, Availability, Confidentiality.

   C. This answer is correct. The correct labeling, according to the Top Threats Analysis Methodology, is: Integrity, Availability, Confidentiality.

   D. This answer is incorrect because Integrity, Availability, Confidentiality is the correct labeling, according to the Top Threats Analysis Methodology.

3. Match the monitoring terms in column A with the monitoring descriptions in column B.

Column A: Monitoring Terms	Column B: Monitoring Descriptions
1. D Continuous	D. With little or no delays
1. C Graphic visualization	C. Facilitates data presentation and analysis
1. B Comprehensive	B. Necessary relevant statistics
1. A Forensic	A. Investigatory friendly

Chapter 5:

Cloud Auditing

Cloud Auditing

Cloud Auditing

5.1 Learning Objectives

After completing this chapter, learners will be able to:

1. Describe the compliance program evaluation approach.
2. Recall the governance perspective.
3. Outline the perspectives of laws, regulations and standards.
4. Define service changes.
5. Explain the need for continuous assurance and continuous appliance.

5.2 Overview

Auditing is not an idea but a process, as defined in ISO 19011. Audit or auditing is often associated, if not almost synonymous, with financial auditing, and it can refer to an independent examination of the financial information of any entity. In the context of this guidance, the focus is on auditing of an IT system, specifically a cloud-based IT system.

Auditing is defined in this guidance as the independent assessment conducted by a qualified assessor of the conformity of the internal and external (cloud) processes within the scope of the applicable regulatory requirements, organizational policies and standard requirements.

This chapter covers aspects related to the following:

- The differences between internal and external cloud auditing
- The differences between auditing and assessment
- The most relevant auditing standards
- Auditors' competency requirements
- The differences between auditing on-premises infrastructure vs. cloud
- Building and executing a cloud audit plan
- How the cloud fits into the existing company approach to auditing
- IT security audit characteristics and criteria (see section 1.4.15).

5.3 Audit Characteristics, Criteria and Principles

This section introduces key security audit characteristics, criteria and principles to provide guidance on what it is to be audited and how an audit should be performed in a cloud environment. In addition, it highlights the differences between internal and external cloud auditing, and between an audit and an assessment. Corresponding to the discussion of audit and assessment, the terms auditor and assessor are used interchangeably in many of the guidance chapters, even though they can represent different roles, responsibilities and customer expectations, depending on the program.

5.3.1 Key IT Security Audit Characteristics and Criteria

This section describes key characteristics and criteria of an IT security audit.

Requester

This is the person or group requesting the audit (**figure 5.1**). The audit team typically defines the key questions the audit should address, the expectations around scope, coverage (breadth vs. depth), the methodology and reporting. In the context of the cloud there may be multiple auditees, and the requester may change depending on which group is performing the audit. The auditor should request relevant information and context from the requester and agree on what level of support, if any, can be provided during the audit process.

Figure 5.1—Auditor's Vantage Point

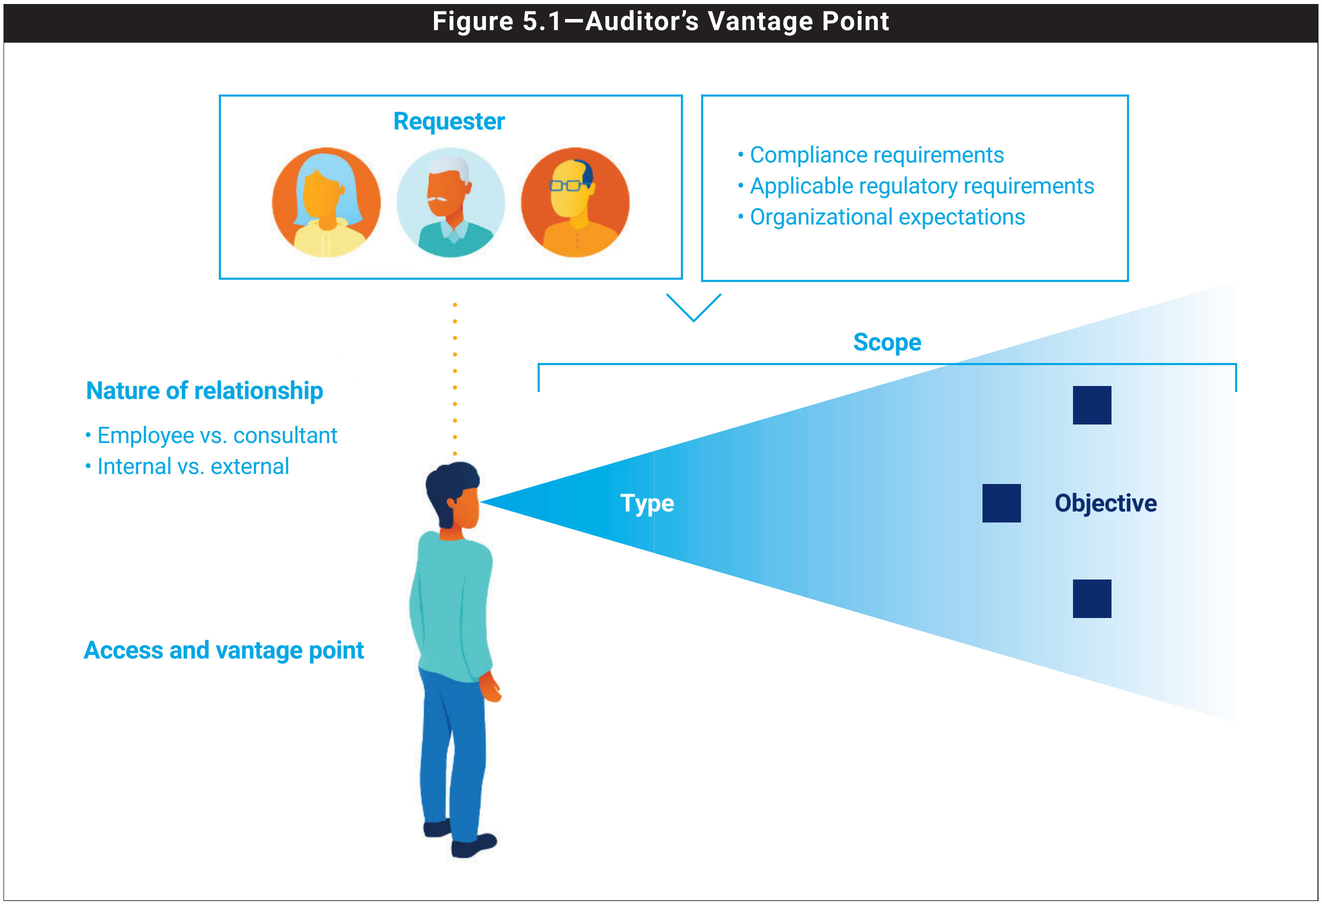
Requester
• Compliance requirements
• Applicable regulatory requirements
• Organizational expectations
Scope
Nature of relationship
• Employee vs. consultant
• Internal vs. external
Type
Objective
Access and vantage point

Scope

The audit scope identifies the systems and processes to be assessed. Prior to starting an audit is important to understand its scope and whether it aligns with the stated objectives and purpose. Each audit should have a clear statement describing its aim, its goals, and the factors driving the need for an audit. From an internal audit perspective, there could be a methodology to assess scope on a functional basis (e.g., an enterprise cloud email audit, a telecom cloud audit) or on a business unit basis (procurement, HR). The scope of an organization "born in the cloud"—one that solely uses cloud services and end-user devices—will be materially different from that of a traditional enterprise engaged in hybrid cloud initiatives. Auditors will need to understand and assess the "vanishing perimeter," as data traffic increases from portable devices that traditionally would have gone through perimeter security connected directly to the cloud. The entity that is the focus of the organization compliance efforts (e.g., PCI SSC Cloud Computing Guidelines) would determine the scope of a compliance audit.

Purpose and Objectives

The key drivers and rationale for the audit should consider these questions: Why is an audit needed? To verify adherence to internal policy? To investigate control failures in relation to persistent breaches or material policy infringements? Some audits may be mandated by a regulator or requested by a customer as a condition of procuring services. In many cases, a CSP seeks an audit to achieve or confirm a certification or attestation.

Type

Audits may be categorized according to these types: IS/IT audit, compliance audit, third-party service audit, forensic audit or financial audit. In addition, they can be categorized as internal (first party) or external (second or third party), and by type of cloud (public, private, hybrid) and service delivery (SaaS, IaaS, PaaS).

Access and Vantage Point

Audit procedures must be discussed and addressed before fieldwork begins, and reflected in an internal audit charter, an external audit engagement letter, or protocols established by regulatory bodies. The discussion should include access or privileges granted to the auditors, and the vantage point from which they will conduct the audit. The scope may include physical access to a facility or a restricted area of a building; logical access (e.g., to a cloud tenant); or controlled network access. In practice, an audit of a virtual data center may be exclusively remote from the physical CSP premises, because even though a data center audit may be requested, some CSPs do not permit customers to perform on-site data center audits. In such cases, a CSP may provide an attestation or audit report from an independent third party as evidence of its control environment.

Nature of the Relationship

This defines the nature of the contractual relationship between auditors, auditees and requester—e.g., whether the auditor is an employee or consultant of the organization that is the subject to the audit. An audit carried out by an insider is commonly referred to as an internal audit. Auditors are bound by principles of independence and operate on the basis of their objectivity. An external audit means the auditor is operating independently of auditee management, e.g., a PCI auditor from an accountancy firm. In general, independent audits are considered more trustworthy and hence more valuable to external parties. In some circles, external audit refers to auditors who opine on an organization’s financial statements. Knowing that information technology is the basis upon which transactions are processed and financial reports generated, the external audit team will also include IT auditors.

Audit Resourcing

The following considerations apply to accessing resources and staffing for the audit:

- Is the audit part of a broader audit with other auditors examining other elements of the organization?
- Is it a solo audit or is a specialized team required to properly cover the scope?
- What level of information sharing and coordination between auditors is appropriate and beneficial?
- Is the audit part-time or full-time?

For internal audits, there is guidance related to competency with respect to audit areas requiring specific skills and experience. Also, in keeping with objectivity, there is guidance related to cooling off periods. For example, an employee in an audit role who previously was engaged in managing an area should not be involved in auditing that area for a certain period of time, due to segregation-of-duties requirements.

Sequencing

Questions to consider regarding the context and space/time location of the audit:

- Is the audit part of a sequence of audits? For example, does it build, relate to, or repeat a prior audit?
- Should an input of this audit be the output of a prior audit? Is the prior scope still relevant and appropriate?
- Should the planned scope be updated to reflect a change in audit objectives, purpose, vantage point or access?

Standards

There are two important categories of standards to consider for use during the auditing process:

- **Standards that define an auditing methodology**—These define how an auditor should conduct an audit and state the rules and requirements that auditors must follow. This category includes, for example, ISO/IEC 17021, 19011, 27006, ISO/IEC 27017 and 27007, or ISAE 3000/ISAE 3402/SOC2, or STAR Certification or STAR Attestation; for privacy review, ISO/IEC 27701 and ISO/IEC 27018. While not specifically related to IT audit, the Institute of Internal Auditors (IIA) provides auditing methodology guidance in its International Professional Practices Framework (IPPF). ISACA provides its audit methodology guidance in its Information Technology Assurance Framework (ITAF).
- **Technical standards that define the control specification against which a process/service/product is to be evaluated**—That is, the security control framework introduced in section 2.5. These can be compared against the control claims made by the auditee to identify gaps and differences, which the auditor needs to evaluate in the context of the audit purpose and objectives and a statement of applicability.

Constraints

This defines factors or elements that may restrict or limit audits. It is vital to proactively identify constraints during audit planning, because they may have a downstream impact and negatively affect an audit. For example, time constraints may limit the scope, audit methodology, and the number of key controls that can be checked. An auditee's operational policy may prohibit certain types of activities or access by third parties. It may be possible to negotiate easing or removal of constraints, particularly where they unreasonably impact the audit or auditor. Again, this is best done prior to the audit commencing. In some cases, if constraints are too inhibiting, the audit may have to be delayed until a mutual solution can be agreed upon. Some important questions to consider when considering the constraints: How might they help or confine the scope? How do the constraints support or impact the audit methodology, process or timing? Are there any constraints that materially undermine the feasibility of the audit (e.g., gaining access to subject matter experts or facilities)? If an organization has invested in material capability improvements, is the prior audit standard still applicable, or should a stricter standard apply? If the same audit standard is used, should a different level of assurance apply (e.g., desktop review vs. sampling log files)? The opinion or audit report should include any constraints or exceptions to the report due to the constraints.

Reporting

Reporting is the communication that takes place during the life cycle of an audit. In general, the reporting follows the various phases of an audit and includes communication of the intent to conduct an audit, (unless it is a surprise audit); interim reports during the course of the audit (especially for lengthy audits); and a final report at the completion of the audit that includes findings and recommendations. If the audit report requests a response in the form of a corrective action plan (CAP), a post-audit report may be created to provide confirmation of the receipt of the CAP and any comments regarding its content.

Consideration may be given to communication during other audit phases: kickoff meeting (who attended and what was discussed?); fieldwork status (e.g., if a significant finding is noted, what is the communication protocol—to whom, by whom, in what format?). Communication continues after the audit is finalized. For instance, the internal auditors will need to work with management to resolve audit findings. In addition, the auditor, typically the audit manager or the audit department, also communicates the status of audit findings to the executive team or the audit committee of the board of directors.

The audit report has a defined audience, such as internal decision makers, external stakeholders or regulators. The findings from one audit report may form part of, or be referenced in, a broader organizational audit.

5.3.2 Auditing Core Principles

There are global and country-based professional standards bodies that represent both internal and external auditors. Each has developed a series of professional standards and practices to guide its members. They are designed to offer a set of core principles to guide behaviors and decision-making among practitioners and workers in the audit profession.[219] Following are some of the most frequently mentioned principles for auditors**.**

Independence

Independence is essential to the effectiveness of the audit function. The auditor shall be free from any bias and conflict of interest. Independence is one of many principles governing audits as described in ISO 19011. The principle of independence shall be adhered to not only in mind but also in appearance. Also, the internal auditor shall resist any undue pressure or interference in establishing the scope of the assignments or the manner in which assignments are conducted and reported, in case any should deviate from set objectives. Closely related to the principles of both independence and due professional care is the concept of professional skepticism, defined as an attitude that includes a questioning mind and a critical assessment of audit evidence.

Integrity and Objectivity

The auditor shall be a person of high integrity, honest and truthful. This individual shall operate in a highly professional manner and be viewed as fair during all phases of the audit. The individual shall avoid all conflicts of interest and not seek to derive any undue professional benefit or advantage from the audit role.

Due Professional Care

The auditor shall pay particular attention to certain key audit objectives, such as establishing the scope of the engagement to prevent the omission of important aspects, recognizing the risks and materiality of the areas, having required knowledge and skills to review complex matters, and establishing the extent of testing required to achieve the objectives within specified deadlines.

Confidentiality

The auditor shall keep confidential information secure from others during all the audit phases until the information can be destroyed after being securely archived for the required period of time. (The time periods vary according to the record-keeping requirements set for the auditor.) Confidential information shall not be shared with people outside the engagement team or with third parties outside the company without specific approval of the client, unless there is a legal or professional responsibility to do so.

[219] McCafferty, J.; "Ten Core Principles for Internal Auditors to Live By," Internal Audit 360°, 29 November 2018, https://internalaudit360.com/ten-core-principles-for-internal-auditors-to-live-by/2/

Skills and Competence

An auditor who lacks certain expertise shall procure the required skills by enlisting in-house experts or engaging the services of an outside expert, provided that doing so does not compromise the auditor's independence. The objective is to ensure that the audit team as a whole has all the expertise and knowledge required for the area under review.

Risk-Based Audit

The auditor will perform a risk assessment exercise to identify the important audit areas and to tailor the audit activities so that the detailed audit process is prioritized in favor of high-risk areas and issues, while less time is devoted to low-risk areas due to curtailed audit procedures. Additionally, this approach shall ensure that risks under consideration are more aligned to the overall strategic and company objectives rather than narrowly focused on process objectives.

System and Process Focus

An auditor should adopt a system- and process-focused methodology in conducting audit procedures. This methodology is more sustainable than the one adopted to test transactions and balances, because it goes beyond error detection to include error prevention. It requires a root cause analysis to be conducted on deviations to identify opportunities for system improvement or automation with the aim of strengthening the process and preventing repetition of such errors.

Avoiding Participation in Operational Decision-Making

The auditor shall avoid participation in operational decision-making, which may be subject to a subsequent audit. Such activity would violate ethics concerning objectivity and bring into question the validity of the audit and its findings.

Sensitivity to Multiple Stakeholder Interests

The auditor shall evaluate the implications of audit-related observations and recommendations on multiple stakeholders, especially regarding diverse interests that may be conflicting in nature. In such situations, the auditor shall remain objective and present a balanced view. This approach permits senior management to make a decision using all the available information, and to balance the strategic objectives of the company with the expectations and interests of its multiple stakeholders.

Quality and Continuous Improvement

Along with continued professional development, auditors should seek quality assessments from external resources to identify improvement opportunities. Because no process or function is ever perfect, internal audit should constantly be looking for ways to improve.

Auditors

Around the world many cases are guided by standard-setting bodies, primarily the International Standards Organization (ISO), the International Federation of Accountants, and the International Auditing and Assurance Standards Board.

5.3.3 Internal vs. External Auditing

Audits that can be broadly categorized in two groups: internal and external audit. The differences between them are presented in the following sections.

Internal Audit

Internal audit is an independent, objective assurance and advisory activity designed to add value and improve an organization's operations. It helps an organization accomplish its objectives through use of a systematic, disciplined approach to evaluate and improve the effectiveness of risk management, control and governance processes.[220] The audit function should report to the top management in planning, providing status updates and reporting on all audit results.

Internal audits typically have a broader scope than external audits because internal audits aim to drive improvements in compliance and in the overall business. Therefore, the internal audit scope may include certifications and regulatory requirements, as well as applicable internal business requirements. The internal auditors must identify where risk exists and make recommendations to adequately mitigate such risks. In many cases, internal audits may cover areas not addressed or considered out of scope by external auditors. It is recommended that internal audits be carried out at planned intervals. It is also advisable to create a plan that covers the entire year to ensure that all critical aspects of the business and management system are covered year over year. The audit function can (and should) deviate from the annual plan if audit resources are needed to address a new area, risk, implementation, etc.

External Auditing

External auditing represents a distinct segment in the professional and accredited third-party services market. An external auditor performs an audit in accordance with specific laws, rules, standards and regulations, and is independent of the entity being audited.[221, 222, 223]

External audits can be classified in two broad categories:

- **Third-party audits**—External audits performed in accordance with internationally accepted practices such as ISO/IEC 27001, in conjunction with ISO/IEC 19011, and ISAE 3000. These audits are typically linked to a certification or attestation process.
- **Second-party audits**—Audits performed to satisfy a contractual obligation for a customer, or to fulfill a regulatory requirement or for marketing purposes. They often follow government or industry regulatory requirements. The external auditor may check the work of internal auditors as part of the process to reassure the organization on the reliability of an internal control.

The services of external auditors are primarily devoted to examining compliance with applicable controls and rendering of opinions on their effectiveness. Unlike internal auditors, external auditors are prohibited from advising on risk mitigation. They are allowed to share general best practices or provide professional authoritative references that an organization can review, which may facilitate its mitigation efforts.

Internal and external auditors are both important for every organization, to assess its overall compliance and governance processes and the effectiveness of the applicable controls it implements. Both are responsible for

[220] The Institute of Internal Auditors, "Definition of Internal Auditing," https://global.theiia.org/standards-guidance/mandatory-guidance/Pages/Definition-of-Internal-Auditing.aspx

[221] *Op cit* ISO, ISO/IEC 27001 *Information Security Management*

[222] ISO, ISO 19011:2018 *Guidelines for auditing management systems*, July 2018, www.iso.org/standard/70017.html

[223] IAASB, "International Standard on Assurance Engagements (ISAE) 3000 Revised, Assurance Engagements Other Than Audits or Reviews of Historical Financial Information," 9 December 2013, www.iaasb.org/publications/international-standard-assurance-engagements-isae-3000-revised-assurance-engagements-other-audits-or-0

following professional standards in their dealings with clients. Internal and external audit play a vital role in driving the organization's improvement, and in giving a true and fair view of the organization overall. The role of the external auditor is critical, because it certifies the integrity of the organization's processes. See **figure 5.2** for an example of some major characteristics of internal vs. external audit.

Figure 5.2—Basis of Comparison Between Internal and External Audits

	Internal Audit	External Audit
Meaning	● Ongoing audit function performed within the organization by an internal auditing team	● Audit function performed by independent firm
Conducted by	● Carried out by employees/contractors of the organization, appointed by management	● Auditor from third party, appointed by the regulator or the accredited, certifying organization
Reports	● Management	● Management
Audit period	● Continuous process at planned intervals	● Consistent cycle as specified by regulation or standard
Purpose	● To evaluate routine compliance activities and provide control for improvement	● To investigate and verify compliance and effectiveness of the organization's management and/or business systems
Legally bound	● No, internal audit is not compulsory	● Compulsory per government act or international standards
Scope	● Decided by management	● Decided by government body or per rules and regulations associated with standards

Source: EDUCBA, "Internal Audit Vs External Audit," www.educba.com/internal-audit-vs-external-audit/

5.3.4 Audit and Assessment

In a business context, it is very common to encounter the terms audit and assessment, sometimes used interchangeably. For example, it would be logical to prepare for a security audit by conducting a security assessment beforehand. The assessment report would typically identify instances of noncompliance with whatever standard of measure was being used, and recommendations for activities or actions to remediate any deficiencies or gaps identified. By contrast, an audit is intended to assess how well an organization is meeting a set of external standards, laws or regulations, and to validate an organization's alignment with its own internal standards and policies.

As illustrated in **figure 5.3**,[224] there are various activities that auditors can carry out, some of which are specific to their profession.

Although the types of services that internal and external auditors provide can be similar in nature, there are some specific differences in objectives set or techniques applied for the processes summarized in **figure 5.3**.

[224] Louwers, T.; A. Blay; D. SInason; J. Strawser; J. Thibodeau; *Auditing & Assurance Services*, 7th Edition, McGraw-Hill Education, US, 2017

Figure 5.3—Select Categories of Auditing/Evaluation Activities

Category	Description
Assurance	The objective of an assurance engagement is for members to evaluate or measure a subject matter that is the responsibility of another party against identified suitable criteria and to express a conclusion that provides the intended user with a level of assurance about that subject matter. Assurance engagements are intended to enhance the credibility of information about a subject matter by evaluating whether the subject matter conforms in all material respects with suitable criteria, thereby improving the likelihood that the information will meet an intended user's needs. In this regard, the level of assurance provided by the conclusion conveys the degree of confidence that the intended user may place in the credibility of the subject matter.[225]
Auditing	Auditing is a type of assurance service that includes a systematic process of objectively obtaining and evaluating evidence regarding assertions about certain subject matters. Auditing ascertains the degree of correspondence between the assertions and established criteria and involves communicating the results to interested users.[226]
Attestation	Attestation engagements are another category of assurance services in which an external auditor is engaged to issue a report on subject matter that is the responsibility of another party, e.g., agreed-upon procedures, reporting on controls at a service organization.[227]
Assessment	Assessment goes further than an audit, as it involves the determination of actions necessary to make the assessed entity compliant.[228]
Evaluation	Evaluation may be considered a process of delineating, obtaining and providing useful information for judging decision alternatives. Evaluation is the determination of the merit or worth of something.[229]

5.4 Auditing Standards for Cloud Computing

This section presents the auditing standards and auditor competencies for cloud computing.

5.4.1 Auditing Standards

This section describes the standards for auditing cloud computing.

ISO/IEC Standards

The ISO/IEC standards include a series of international standards, codes of practice, and requirements for auditing management systems in general, and information security management systems in particular.

- ISO/IEC 17021-1:2015 *Conformity assessment—Requirements for bodies providing audit and certification of management systems—Part 1: Requirements*
  - Establishes principles and requirements for bodies providing audit and certification of management systems, such as impartiality, competence, responsibility, openness, responsiveness to complaints, general requirements, structural requirements, resource requirements, information requirements, process requirements, and management system requirements for the auditing body.
- ISO/IEC 27006:2015 *Information technology—Security techniques—Requirements for bodies providing audit and certification of information security management systems*
  - Specifies requirements and provides guidance for bodies providing audit and certification of an information security management system (ISMS) in addition to the requirements contained within ISO/IEC 17021-1 and

[225] IAASB, "Assurance Engagements – Completed," December 2003, www.iaasb.org/projects/assurance-engagements-completed

[226] American Accounting Association, *A Statement of Basic Auditing Concepts*, USA

[227] AICPA, "Statement on Standards for Attestation Engagements No. 18," April 2016, www.aicpa.org/content/dam/aicpa/research/standards/auditattest/downloadabledocuments/ssae-no-18.pdf

[228] Transition Support, ISO 9001 FAQs, https://transition-support.com/faqs.htm

[229] Lawrenz, F.; M. Thao; "Curriculum Evaluation," *Encyclopedia of Science Education*, 6 March 2014, https://link.springer.com/referenceworkentry/10.1007%2F978-94-007-6165-0_150-1

    ISO/IEC 27001. It is primarily intended to support the accreditation of certification bodies providing ISMS certification. The requirements contained in this international standard need to be demonstrated in terms of competence and reliability by any body that provides ISMS certification. The guidance contained in this international standard provides additional interpretation of these requirements for any body providing ISMS certification. This international standard can be used as a criteria document for accreditation, peer assessment, or other audit processes.

- ISO 19011:2018 *Guidelines for auditing management systems*
    - Provides guidance on auditing management systems, including the principles of auditing, managing an audit program, and conducting management system audits. It provides guidance on evaluating the competence of individuals involved in the audit process, including the individuals managing the audit program, auditors and audit teams. It is applicable to all organizations that need to plan and conduct internal or external audits of management systems, or manage an audit program. The application of this document to other types of audits is possible, provided that special consideration is given to the specific competence needed.
- ISO/IEC 27007:2020 *Information security, cybersecurity, and privacy protection—Guidelines for information security management systems auditing*
    - Provides guidance on managing an information security management system (ISMS) audit program, on conducting audits, and on evaluating the competence of ISMS auditors, in addition to the guidance contained in ISO 19011. This document is applicable to those needing to understand or conduct internal or external audits of an ISMS or to manage an ISMS audit program.
- ISO/IEC 27701 *Security techniques—Extension to ISO/IEC 27001 and ISO/IEC 27002 for privacy information management—Requirements and guidelines*
    - Specifies requirements and provides guidance for establishing, implementing, maintaining and continually improving a privacy information management system (PIMS) in the form of an extension to ISO/IEC 27001 and ISO/IEC 27002 for privacy management within the context of the organization.
- ISO/IEC 27018 *Information technology—Security techniques—Code of practice for protection of personally identifiable information (PII) in public clouds acting as PII processors*
    - Establishes commonly accepted control objectives, controls and guidelines for implementing measures to protect personally identifiable information (PII) in line with the privacy principles in ISO/IEC 29100 for the public cloud computing environment.
- ISO/IEC 27017:2015 *Information technology—Security techniques—Code of practice for information security controls based on ISO/IEC 27002 for cloud services*
    - Provides guidelines for information security controls applicable to the provision and use of cloud services through the following:
        - Additional implementation guidance for relevant controls specified in ISO/IEC 27002
        - Additional controls with implementation guidance that specifically relate to cloud services
    - This international standard provides controls and implementation guidance for both cloud service providers and cloud service customers.

International Standards for the Professional Practice of Internal Audit (IIA)

The IIA International Standards for the Professional Practice of Internal Auditing (referred to as the standards in this section) are essential in meeting the responsibilities of internal auditors and the internal audit activity. If laws or regulations prohibit internal auditors or the internal audit activity from conformance with certain parts of the standards, conformance with all other parts of the standards and appropriate disclosures is needed.

If the standards are used in conjunction with frameworks issued by other authoritative bodies, internal audit communications may cite the use of other frameworks as appropriate. In such cases, if inconsistencies exist between the standards and other frameworks, internal auditors and the internal audit activity must conform with the standards and may conform with the other frameworks if they are more restrictive.

Purposes of the standards:

- Delineate basic principles that represent the practice of internal auditing
- Provide a framework for performing and promoting a broad range of value-added internal auditing
- Establish the basis for evaluation of internal audit performance
- Foster improved organizational processes and operations

The standards are principles-focused mandatory requirements:

- Statements that set forth basic requirements for the professional practice of internal auditing and for evaluating performance effectiveness, which are internationally applicable at organizational and individual levels
- Interpretations, which clarify terms or concepts within the statements[230]

ISAE Standards

Assurance Engagements Other than Audits or Reviews of Historical Financial Information (ISAE 3000) describes general requirements for the qualification and conduct of an auditor (e.g., professional judgment and skepticism) and for accepting, planning and carrying out an audit engagement. It is a high-level auditing standard that provides the required high-level framework. Issued by the International Auditing and Assurance Standards Board (IAASB), the standard includes general requirements for audit criteria without specifying details. The standard distinguishes between audits with reasonable assurance and audits with limited assurance and distinguishes so-called attestation engagements from so-called direct engagements.

ISAE 3000 addresses several important issues relating to practitioners (auditors) and subjects of assurance (auditees) that are the code of conduct for external assurance on nonfinancial information. ISAE 3000 is supported by the International Framework Ethics for Assurance Engagements (the IAASB Framework), which defines professional accountants and describes the elements and objectives of an assurance engagement and auditor firm. The standard covers all aspects of an assurance engagement, including engagement acceptance, agreement to the terms of engagement, planning and performing the engagement, using the work of experts, obtaining evidence, considering subsequent events, documentation, and preparing the external assurance report.

Assurance Reports on Controls at a Service Organization (ISAE 3402) was developed to provide an international assurance standard for allowing certified public accountants (CPAs) to issue a report for use by customer organizations and their auditors (customer auditors) on the controls at a cloud service provider likely to impact or be a part of the customer organization's system of internal control over financial reporting. The auditor shall not represent compliance with ISAE 3402 unless the auditor has complied with the requirements of ISAE 3000 as well.

Like ISAE 3000, ISAE 3402 states that an assurance engagement may be a reasonable assurance engagement or a limited assurance engagement; that an assurance engagement may be either an "assertion-based" engagement or a "direct reporting" engagement; and that the assurance conclusion for an assertion-based engagement can be worded either in terms of the responsible party's assertion or directly in terms of the subject matter and the criteria. ISAE 3402, however, deals only with assertion-based engagements that convey reasonable assurance, with the assurance conclusion worded directly in terms of the subject matter and the criteria. It applies only when the cloud service provider is responsible for, or otherwise able to make an assertion about, the suitable design of controls.

[230] The Institute of Internal Auditors, "International Standards for the Professional Practice of Internal Auditing (Standards)," October 2008, revised October 2012, https://na.theiia.org/standards-guidance/Public%20Documents/IPPF%202013%20English.pdf

The standard defines the guidelines concerning the following themes:

- Ethical requirements
- Management and those charged with governance
- Acceptance and continuance
- Assessing the suitability of the criteria
- Materiality
- Obtaining an understanding of the service organization's system
- Obtaining evidence regarding the description
- Obtaining evidence regarding design of controls
- Obtaining evidence regarding operating effectiveness of controls
- The work of an internal audit function
- Subsequent events
- Documentation
- Preparing the service auditor's assurance report.

SSAE 18

SSAE, or Statement on Standards for Attestation Engagements, is overseen by the American Institute of Certified Public Accountants (AICPA) and specifically, the Auditing Standards Board (ASB).[231]

According to the AICPA, "Service Organization Control (SOC) reports are internal control reports on the services provided by a service organization providing valuable information that users need to assess and address the risks associated with an outsourced service." SSAE is used to standardize the way organizations document their implementation and operation of various compliance controls, and how they report on compliance controls, including SOC1, SOC 2, and SOC 3 reports.

Under the SSAE 18 guidelines (revised in 2016), service organizations need to have specific management programs for their third-party vendors. If an organization has third-party vendors, also known as subservice organizations, the company needs to have clearly described responsibilities for each of these vendors. In addition, it needs to have recorded performance reviews that contain routine audits and reviews on what it learned from those findings.

Service organizations also need to have a formal process to gauge the efficiency of the risk assessment. This new statement of standards addresses risks and mandates an assessment for them. As a part of each report for third-party vendors, each company needs to include specific plan details on how it addresses risk management. The report for this program also needs to explain and outline the efficiency of this plan.

Most importantly, SSAE 18 is derived from ISAE 3000—the AICPA is a close partner with the International Accounting and Assurance Standards Board (IAASB).

5.4.2 Auditor Competency

Auditing the security and privacy of cloud computing, regardless of service or deployment model, requires a combination of a variety of skill sets, expertise and experience.

They can be divided into three families of skills:

[231] AICPA, "Clarified Statements on Standards for Attestation Engagements," https://www.aicpa.org/research/standards/auditattest/ssae.html

- Auditing skills and expertise
- Cloud security skills and expertise
- Cloud technology skills and expertise

Under the first category, auditing skills and expertise, there are essentially two main groups of standards, best practices and approaches, with related professional certifications:

- ISO professional accreditation/certification
- National institutes of chartered accountants and IIA accreditation/certification

Auditing Skills and Expertise

ISO professional accreditation/certification—Certifications based on the International Standards Organization standards consist of the following:

- **ISO/IEC 27001 Lead Auditor certification**—Consists of a professional certification for auditors specializing in information security management systems (ISMS) based on the ISO/IEC 27001 standard and ISO/IEC 19011.
- **Certified CSA STAR Auditor**—For ISO/IEC 27001 qualified auditors working for certification bodies accredited by an International Accreditation Forum (IAF) member to ISO/IEC 27006 and approved to carry out CSA STAR Certification audits. The objectives are to teach how to carry out a STAR Certification audit to assess compliance with the Cloud Security Alliance (CSA) Cloud Controls Matrix (CCM) and to teach how to carry out a maturity assessment, determine a maturity score, and recommend a rating for STAR Certification.
- **National institutes of accountant professional accreditation/certification**—Standards, guidelines and professional certifications issued by the International Auditing and Assurance Standards Board (IAASB) and the various national institutes of accountants (AICPA in the US, ICAEW in England and Wales, IDW in Germany, CNDCEC in Italy, ICAI in India, etc.)
- **Certified Public Accountant (CPA)**—The CPA designation identifies licensed accounting professionals who offer financial statement audits and other attestation or advisory services.[232]
- **Chartered Accountants (CA)**—A chartered accountant (CA) is an international accounting designation granted to accounting professionals in many countries around the world, aside from the United States. In the US, the equivalent of the CA is the CPA designation.[233]

Example of other relevant professional accreditations/certifications:

- **Certified Information Systems Auditor (CISA)**—CISA is a global certification managed and administered by ISACA. It is designed for those who audit, control, monitor and assess an enterprise's information technology and business systems.[234]
- **Certified Internal Auditor (CIA)**—is a certification offered to accountants who conduct internal audits. The Certified Internal Auditor designation is conferred by the Institute of Internal Auditors (IIA) as the only globally recognized internal audit certification.[235]
- **Certified Fraud Examiner (CFE)**—The Association of Certified Fraud Examiners is a professional organization of fraud examiners; it administers and grants the CFE.[236]

Cloud Security Skills and Expertise

Note: A cloud computing security auditor needs knowledge of cloud security in general and possibly a good understanding of the specific technology and platform under evaluation.

[232] Kenton, W.; "What Is a Certified Public Accountant (CPA)?," Investopedia, 21 August 2019, www.investopedia.com/terms/c/cpa.asp
[233] Top Accounting Degrees, "What is a Chartered Accountant?," www.topaccountingdegrees.org/faq/what-is-a-chartered-accountant/
[234] ISACA, "CISA," www.isaca.org/credentialing/cisa
[235] The Institute of Internal Auditors, "Prove Credibility & Proficiency," https://na.theiia.org/certification/CIA-Certification/Pages/CIA-Certification.aspx
[236] ACFE, "CFE Qualifications," www.acfe.com/cfe-qualifications.aspx

In terms of the first category, knowledge of cloud security in general, the following certifications are appropriate:

- **CSA Certificate of Cloud Security Knowledge (CCSK)**—The CCSK, which was the first exam to test an individual's understanding of a full range of cloud security topics, assesses an individual's understanding of security strategies that are necessary in today's cloud environments. Now in its fourth iteration, The CCSK is predominantly based on CSA's "Security Guidance for Critical Areas of Focus in Cloud Computing." The self-proctored, open book exam poses 60 questions over 90 minutes to assess the individual's understanding of cloud security. CCSK is a combination of true/false and multiple choice questions. It is available in English, Spanish and Japanese.
- **ISC2 Certified Cloud Security Professional (CCSP)**—The CCSK, initially a collaboration between CSA and ISC2, brings the tactical considerations of cloud security into focus as it tests essentially the same security domains as the CCSK. These tactical considerations require the individual to understand how the implementation of controls would be carried out in an on-the-job scenario. CCSP, currently in its second iteration, is administered in a Pearson VUE Testing Center. It consists of 125 questions to be answered in 180 minutes. CCSP is multiple choice questions only and is available in English.
- **GIAC Cloud Security Automation (GCSA)**—The SANS Institute's certification entity GIAC offers the GCSA, which certifies individuals in applying Secure DevOps principles, practices and tools for building and deploying secure software and infrastructure in the cloud. The exam entails the completion of 75 questions within a two-hour time limit. The minimum passing score is 61%. This exam is web-based and can be completed either via remote proctor by ProctorU or on-site proctor by PearsonVue.

Cloud Technology Skills and Expertise

Besides having knowledge and experience in applying relevant auditing standards (from ISO, IAASB, etc.), a cloud auditor needs to have dedicated knowledge and working experience with the specific technology and business sector in the scope of assessment.

There are several product training programs and certificates from vendors such as Amazon, Microsoft, Google, Oracle and IBM. All of these training programs and accreditations contain gaps, both in the depth and breadth of coverage with respect to cloud security assurance, which is why CSA built the CCAK.

5.5 Auditing an On-Premises Environment vs. Cloud

This section discusses the levels of assurance that are required to satisfy cloud customer requirements and the differences between auditing on-premises and off-premises cloud systems.

5.5.1 Understanding the Required Level of Assurance to Satisfy the Cloud Customer

It is important for organizations that adopt cloud computing technologies to revise and enhance their cloud governance programs to accommodate the impact of cloud technologies on their business processes and strategies. This normally includes the governance strategy and selection of a relevant framework; modification to current risk management programs to address some unique risk factors introduced by cloud computing; revisions and additions to corporate policy necessitated by cloud service and delivery models; the creation of new roles and responsibilities to manage and monitor the new controls and processes introduced by cloud computing; and revisions to legacy information security programs, including existing cybersecurity programs. The review and evaluation of the adequacy and effectiveness of the changes would normally constitute the scope of various audits that would be conducted by the audit function to achieve a commensurate level of assurance to the cloud customer.

Based on the penetration of cloud services into the typical organization technology environment and the two-dimensional aspect of both the number of service providers and the number of organizational instances, it is information security programs, including existing cybersecurity programs. The review and evaluation of the adequacy and effectiveness of the changes would normally constitute the scope of various audits that would be conducted by the audit function to achieve a commensurate level of assurance to the cloud customer.

Based on the penetration of cloud services into the typical organization technology environment and the two-dimensional aspect of both the number of service providers and the number of organizational instances, it is imperative that an enterprise risk management program be implemented or extended to include all known and unknown (shadow IT) cloud services and providers. This risk-based approach will facilitate the creation of a portfolio view of all CSPs, with risk ratings for each type of service provided to the organization. In addition, the risks should be categorized by business objectives such as strategy, operations, compliance and business performance.

5.5.2 Differences Between Auditing On-Premises and Off-Premises Cloud Systems

This section describes the differences between auditing on-premises and off-premises cloud systems.

Technical Considerations

Figure 5.4 provides a high-level overview of the technological differences between cloud computing and on-premises computing. Gaining an understanding of those differences helps in understanding the impact on the auditing.

Figure 5.4—Technology Differences Between Cloud Computing and On-Premises Computing

Consideration	On-Premises Computing System	Off-Premises Cloud Computing
Ownership and management of technology resources	IT assets owned/leased by consuming entity (CE)	IT assets owned/leased by third party/parties
Location of IT resources	IT assets/resources located in facilities owned/leased by CE	IT assets/resources located in facilities owned/leased by third party/parties
Applications and supporting services including infrastructure	Owned and operated by CE in CE facilities	Operated by a third party in third- or fourth-party facility
IT personnel	Employees, contractors and consultants located in CE facilities or remotely with access	IT personnel divided between CE at CE facilities, the third party at third-party facilities, and the fourth party at fourth-party facilities or remotely with access, depending on the cloud service model
Disaster recovery and contingency planning	Performed by CE employees, contractors, and consultants utilizing CE or third-party facilities	Performed by one or more CSPs in accordance with contract terms and conditions (service level agreements)
Application performance, internal controls, security/privacy, and compliance and configurations	Application inventory and related services subjected to enterprise risk assessment by both internal and external audits—External and internal audit develop separate audit plans to fulfill separate objectives. For applications managed/operated by third parties on site, internal audit may conduct site audits subject to contract terms and conditions. For other applications, managed/operated off site, internal audit may conduct off-site audits subject to terms and conditions in contracts.	Inventory of all CSPs, including all services provided subjected to enterprise risk—Scope of audits considers deployments of private and hybrid clouds, and all public cloud services. External audit conducts risk and materiality-based audits of internal controls, security, privacy and compliance to opine on financial statements. Internal audit plans risk-based audits of application services based on the shared responsibility model of CSPs.

Figure 5.4—Technology Differences Between Cloud Computing and On-Premises Computing *(cont.)*

Consideration	On-Premises Computing System	Off-Premises Cloud Computing
Infrastructure/platform audits	Inventory of infrastructure/platform services and products subjected to external and internal audit risk assessments to develop relevant audit plans and approaches	Inventory of all CSPs, including all services provided, subjected to enterprise risk Scope of audits considers deployments of private and hybrid clouds, and public cloud services. External audit conducts risk and materiality-based audit of internal controls, security, privacy and compliance to opine on financial statements. Internal audit plans risk-based audits of application services based upon the shared responsibility model of CSPs.
Governance, risk management and compliance	Review and evaluation of organization programs related to establishment and enforcement of policy and decision making; enterprise risk management and compliance	Review and evaluation of organization programs to manage and monitor cloud provider performance, risk assessment, internal control, compliance responsibilities and contractual requirements
Delivery strategy and architecture	Not applicable	Review and evaluation of an organization's strategy for achieving business objectives through the design and operation of cloud portfolios of services

Differences in Key Aspects of Audit

Figure 5.5 provides a summary of the differences between cloud and on-premises IT in some key aspects of the auditing execution:

- Planning
- Timing
- Approach
- Audit execution
- Reporting and monitoring.

Figure 5.5—Differences in Key Aspects of Audit

Traditional On-Premises IT	Audit of Cloud Computing System
Planning	
Organization-owned and managed technology assets and services are subject to an annual enterprise risk-management assessment. As a result, the audit plan for the year, consisting of various individual audits, is submitted to the board for approval. Once approved, individual audits are put on the calendar based upon various considerations and constraints. The audit schedule, "surprise audits" notwithstanding, is communicated appropriately within the organization and, if relevant, reviewed with regulators. Each on-premises environment requires an assessment and calibration of its relative risk to the organization, its impact and ultimate effect on organizational operations, performance, compliance with laws and regulations, and reputation in the marketplace. Typically, this will also include a review of previous audits and status of incomplete corrective action plans.	Organizations adopting cloud computing platforms and services need to extend the enterprise risk assessment to include all known cloud services. The risk assessment needs to include consideration of the existence and potential impact of shadow IT[237] which is unique to the cloud environment. Based on the number of cloud providers and their significance to the performance and operation of the organization, each service provider instance of operation should be risk-assessed and rated. The final portfolio of risk rating from high to low risk should be considered in planning the number and specific reviews to be included in the annual audit plan. Mature internal audit organizations that practice continuous monitoring programs may opt to leave some capacity in the audit plan to staff an unplanned audit, as the result of a spike in a service provider's risk rating.

[237] Shadow IT is the use of information technology systems, devices, software, applications, and services without explicit IT department approval.

Figure 5.5—Differences in Key Aspects of Audit *(cont.)*

Traditional On-Premises IT	Audit of Cloud Computing System
Timing	
The timing of an individual audit for on-premises facilities may be subject to a number internal considerations and should be reconciled with all affected stakeholders, but it should not conflict with the perceived independence of the auditor. This is important for those situations in which a stakeholder may seek to delay the start of an audit without appropriate justification.	The start time for cloud audits should be up to the organization's internal requirements and influenced by the results of the risk assessment. More importantly, the organization should have negotiated the right to conduct audits, including scope and third-party examinations, in the cloud service agreement.
Approach	
The normal approach for an audit of an on-premises infrastructure is to collect and analyze data, including trend analysis and results of previous audits, to isolate areas warranting specific focus and attention. This is also important for identification of any unique staff skills, experience or competency. Often, the data acquisition and analysis are performed with automated tools available to the audit team. Increasingly, IA departments have adopted agile auditing techniques that incorporate modifications to the process and education of the audit community. Such education is important to both stakeholders and auditees, because it represents a change to audit processes and even approach. This may also influence the use of additional software and techniques for the audit teams to master. Many organizations will have cumulative experience in auditing the various components of on-premises technology and operations, so they will have audit plan and work program templates they can reuse to foster consistency, reliability and audit efficiency.	Audits of cloud service providers are highly dependent upon a few key factors unique to CSPs: ● The cloud service model ● The cloud delivery model ● The shared responsibility model ● The cloud supply chain ● The cloud configuration features Depending upon the auditor's experience and past development of standard audit programs designed to facilitate the audits of CSPs, an audit may be influenced by the time required to design and create new techniques, including software, as opposed to reusing what was developed in the past. In order to identify the processes and controls to include in the audit scope and objectives, the auditor should review the documentation created for the original risk assessment to understand both the underlying risk of the services being provided and the controls necessary to achieve the organization's business objectives. The auditor will be tasked with two major objectives: **1.** How does the CSP ensure that the key security and other controls are properly operating? **2.** How does the cloud customer validate the integrity and reliability of the CSP's processes and controls? The amount of automation and orchestration that many of the leading CSPs use may be an initial challenge for many internal auditors. Depending on the organization's strategy and plans for adopting cloud services, the audit team may want to invest in a program to design and create the tools and processes necessary to audit CSPs efficiently and effectively. This program should include identifying the tools necessary to document and store all the electronic audit evidence necessary to support the audit planning and execution phases.

Figure 5.5—Differences in Key Aspects of Audit *(cont.)*

Traditional On-Premises IT	Audit of Cloud Computing System
Audit Execution	
The majority of audits of on-premises technology resources require visiting the facility and include an examination and review of the physical security. In addition, prior to the commencement of the audit the auditor will have met with the auditee to collect necessary information and review any historical information from previous audits, including corrective action plans. Regardless of the audit style, traditional or agile, the schedule and key milestones will be reviewed with the auditee and other stakeholders.	Audits of cloud providers in some cases will not involve visits to any facilities if they are owned and operated by the CSP, e.g., AWS, Azure. The approach to this will be based on the shared responsibility model outlined in section 6.5.4. Most information and audit evidence will be collected and managed electronically. The amount of audit work and the techniques to be applied will be a result of the staff skills and experience in auditing cloud service providers, and previous investments made in training, education, tools specific to supporting cloud service provider reviews, and the documentation available from the enterprise risk assessment review. The audit team will want to develop a point of contact with the CSP, especially if the provider has audit and process documentation available, including any third-party attestation reports.
Reporting and Monitoring	
The timing and content of preliminary and final audit reports are subject to the type of audit—traditional or agile. The auditee is responsible for responding to the findings in the audit report following review and approval.	Although the CSP is the subject of the audit, the business is the recipient of the report and must respond accordingly. It is up to the business owner to develop the corrective action plan. Regardless of the audit scope—on premises or cloud—the results should be used to update the enterprise risk register, create corrective action plans, and monitor activity over the audit year. A recommended best practice related to reporting and monitoring is to ensure the service agreement with the CSP includes the right to conduct audits, and the CSP's responsibility to respond to audit results and agree to corrective action plans.

5.6 Differences in Assessing Cloud Services

This section discusses the differences in assessing cloud services.

5.6.1 The Impact of the Shared Responsibility Model

When using cloud services, security responsibilities vary in large part depending on the service delivery model used, while responsibility and accountability for compliance remain entirely with the cloud customer. Generally, compliance responsibility reflects the degree of control a party has over the architecture stack. **Figure 5.10** shows a general shared responsibility model depending on the cloud service models.

The cloud service model used (IaaS, PaaS or SaaS) and the cloud deployment model (private, public, hybrid or community) have an impact on the way the organization can leverage the governance tools (auditing, assessment, policies, contracts) to maintain monitoring and oversight over the compliance of the cloud service portfolio.

The assessor's role is to assess not only a particular service, but also the way the enterprise designed and developed it. Considerations include the following:

- How does the service interplay or interact with other services (whether internal or external/cloud services)?
- What is the inventory of data flows across the service supply chain, and what are the information security requirements and likely volumes of flows?
- How is user management handled, and what is the approval process?

- How are decisions about accounts reflected in an audit log?
- Has a third party checked the controls and their design?
- Is it possible to review the conformance tests to see whether the tool or service meets the owner's control objectives?

For each control, the assessor will need to consider the shared responsibility model when evaluating whether a particular control is present and effective. For example, to assess security awareness, an assessor may review policies and standards looking for evidence that management requires all employees to receive regular awareness training, testing (e.g., phishing simulations) and feedback on their performance. In the context of the cloud, the organization must also consider this control through the lens of the shared responsibility model:

- Does the organization (cloud customer) have an awareness training program in place?
- Has the organization (cloud customer) checked that the CSP has an awareness training program in place?

In more generalized terms, for each control the auditor must evaluate the following:

- How does the organization handle the portion of the control it is directly responsible for?
- How does the organization handle the portion of the control that is under the CSP's responsibility, but that the organization is still accountable for?

Note: Two-dimensional thinking is a common theme that must be grasped and applied correctly to think like a cloud auditor.

Taking this theme deeper, following are some suggestions to help an auditor formulate the questions to ask the cloud customer (auditee):

- Is it reasonable to expect a given control to be mirrored at both the CSP and the customer? In the security awareness example, with a cloud customer as the auditee, both organizations would be expected to operate such a control independently of each other. The control objective and design may be very similar at a high level, although the implementation and operation of the control may differ; for example, the CSP may be focused on reducing social engineering of its help-desk staff, whereas the customer (perhaps a bank) may be training its digital teams to recognize Internet-facilitated fraudsters. In the context of cybersecurity and the shared responsibility model, both the customer and CSP have similar control responsibilities, and both must be operated effectively to defend against social engineering attacks.
- Is it reasonable to expect a given control to be layered across both the CSP and customer? For example, the CSP provides its clients with logical access controls to manage access to cloud storage. This user-facing control is underpinned by both logical and physical access controls the CSP operates, i.e., implemented in the cloud platform. In this example, it is clear the overall responsibility for access control is divided between the CSP and its customer. However, it also follows that although the CSP platform can be secure, operating within the CSP's stated control parameters, the customer can suffer a breach for failing to operate the user-facing controls adequately.
- Is a given control operated by entities beyond the CSP and customer? In complex multi-cloud or -party cloud constellations, the assessor will need to consider how many hands are on the controls to identify the control operators involved. Taking this idea further, the assessor should then consider how many fingers of each hand are on the control, assessing the materiality or weighting of each control operator to help direct and estimate the level of effort and analysis the auditor should apply.

In summary, for each control the assessor should take the following steps:

- Identify the nature and placement of in-scope controls, along with control operators.
- Verify the auditee is operating the control in its organization (within the scope of the audit).
- Verify the auditee is asking the right questions of the CSP and other control operators.

Some cloud providers will permit hands-on security testing of workloads and services operating on their platform, subject to permission and constraints. An Infrastructure as a Service (IaaS) provider may allow hands-on security testing of a virtual machine operated by the audit sponsor or account owner. However, the same testing on the cloud authentication service used by all cloud services would be out of scope.

Audit Planning Implications of the Shared Responsibility Model

The division of controls concepts introduced by the various cloud platform deployments pose an audit strategy consideration that takes on greater importance and necessitates focus as the number of cloud providers deployed by the organization expands. The division of controls requires a division of audit approach and practice. As illustrated in **figure 5.6**, the platforms and services of each cloud provider requires a unique combination of test plans due to the shared responsibilities model. While the customer's design and operation of controls can be reviewed, tested and evaluated directly by the audit team, the same procedures cannot be applied to the cloud providers facilities. Hence, the audit plan must be designed to accommodate the practices or protocols of the CSP. ISO/IEC 27017 *Information technology – Security techniques – Code of practice for information security controls based on ISO/IEC 27002 for cloud services*[238] can be referenced as a model describing the details behind shared responsibility. ISO/IEC 27017 gives guidelines for information security controls applicable to the provision and use of cloud services and provides both CSPs and CSCs with controls and implementation guidance and responsibility.

Figure 5.6—Summary of Test Plans

Cloud Provider	Cloud Provider Controls	Customer Controls
CSP 1	Test Plan 1	Test Plan 1
CSP 2	Test Plan 2	Test Plan 2
CSP n	Test Plan n	Test Plan n

Since each CSP may have its own unique collection of tools and services designed to support CSC reviews and audits, the CSC should review and evaluate the capability of the tools and services. Examples of the types of tools and services available[239]:

- **Configuration controls**—Manage automation and configuration assets
- **Compliance controls**—Support adopting various compliance features
- **Encryption controls**—Support various encryption capabilities
- **Access controls**—Enable identify and access management features
- **Logging and monitoring controls**—Support various auditing/audit trail tools
- **Application security controls**—Enable features related to securing existing applications and developing secure ones

It is important to evaluate these portfolios of tools and deploy as appropriate in audits. The audit team should also evaluate the risk to the cloud service owner if the available tools are not being used.[240]

A recent Gartner Report suggests that over the next five years, at least 99 percent of cloud security failures will be the customer's fault.[241] To the extent this prognostication is credible, the importance of audits or assessments of cloud provider platforms should not be underestimated as an organizational best practice.

[238] ISO, ISO/IEC 27017:2015 *Information technology – Security techniques – Code of practice for information security controls based on ISO/IEC 27002 for cloud services*, www.iso.org/standard/43757.html

[239] Dobran, B.; "30 Cloud Monitoring Tools: The Definitive Guide For 2020," 10 August 2020, phoenixNAP, https://phoenixnap.com/blog/cloud-monitoring-tools

[240] *Ibid.*

[241] Boulton, C.; "Posture management: Cloud security tools rise in wake of breaches," 28 February 2020, www.cio.com/article/3529426/posture-management-cloud-security-tools-rise-in-wake-of-breaches

5.6.2 IaaS

In IaaS, the customer does not manage or control the cloud infrastructure (networking, storage, service, virtualization)—it manages operating systems, middleware, deployed applications, and the security features and controls applied to them. From the compliance perspective, that means the assessors will take these steps:

- Evaluate the effectiveness of the control under the direct responsibility of the cloud customer via desktop exercise, vulnerability assessment, penetration testing, first-party audit.
- Evaluate the duty of care applied by the customer in assessing the CSP.

Contracts with cloud IaaS providers will often preclude first- or second-party audits. (CSPs consider on-premises audits a security risk, given that multiple on-premises audits by many customers would present logistical and security challenges.) Consequently, the assessor will need to evaluate other sources of data the customer used in its due diligence effort, such as third-party certification and attestations, self-assessments, white papers, technical documentation, etc. Depending on the audit standard, it may be possible to release actual results only under a nondisclosure agreement (NDA), which means customers will need to enter into a basic legal agreement before gaining access to attestations for risk assessments or other evaluative purposes.

See chapter 7 for further details on how to audit the security of IaaS, PaaS and SaaS using CSA CCM.

5.6.3 PaaS

Further up the cloud stack, as the security responsibilities of the customer diminish, direct control over compliance decreases. In PaaS, the customer is responsible only for data and applications, and the assessor would need to rely more on assessment and audit information the CSP provides, as opposed to direct evaluation of controls implemented by the cloud customers.

5.6.4 SaaS

In SaaS, the customer is responsible only for the security of data, therefore the assessors will largely focus on the compliance offered by the CSP via third-party reports, certifications and attestations. In SaaS, the customer will focus on limiting effective configuration of the client entitlements and ensuring that the SaaS service is used only for certain predefined categories of data.

Very often SaaS providers are smaller organizations, compared to IaaS or PaaS providers, and even compared to some cloud customers. This can sometimes make them more flexible concerning their right-to-audit policy. They might negotiate with the cloud customer the possibility of performing first-party audit and penetration testing.

When the supply chain consists of multiple layers (e.g., SaaS CSP is using a different IaaS CSP to deliver the service), each layer interacts (at least legally) only with the adjacent entity. The assessor should ensure that all possible compliance issues inherent to the CSP subcontractors are under control.

5.6.5 Compliance Inheritance

Many cloud services are certified for various regulations and industry requirements, such as STAR Certification, STAR Attestation, ISO 27001, PCI DSS, SOC 2, HIPAA, and global or regional regulations, like the EU GDPR. These certifications, attestations and authorizations represent an attribute of the service that can be transferred to other services that build on top of the certified ones. For instance, building a SaaS application on top of a STAR-certified IaaS would give the SaaS provider the assurance that its application was created over an infrastructure compliant with the requirements of ISO 27001 and CSA CCM.

From the assessor perspective, there are important factors to consider in order to evaluate the quality of the compliance inherited by a third party (**figure 5.7**):

- Does the cloud customer understand where the boundaries of the compliance inheritance end? Is it aware that it is still ultimately responsible for maintaining the compliance of what it builds and manages? For example, if an IaaS provider is PCI DSS-certified, the customer can build its own PCI-compliant service on that platform and the provider's infrastructure and operations should be outside the customer's assessment scope. However, the customer can just as easily run afoul of PCI and its assessment if does not properly design its own application running in the cloud.
- Is the scope of the CSP certification or attestation relevant? Sometime there is a risk of misinterpreting the real scope of the certification or attestation, especially if the CSP is not transparent about it.

Figure 5.7—Compliance Inheritance

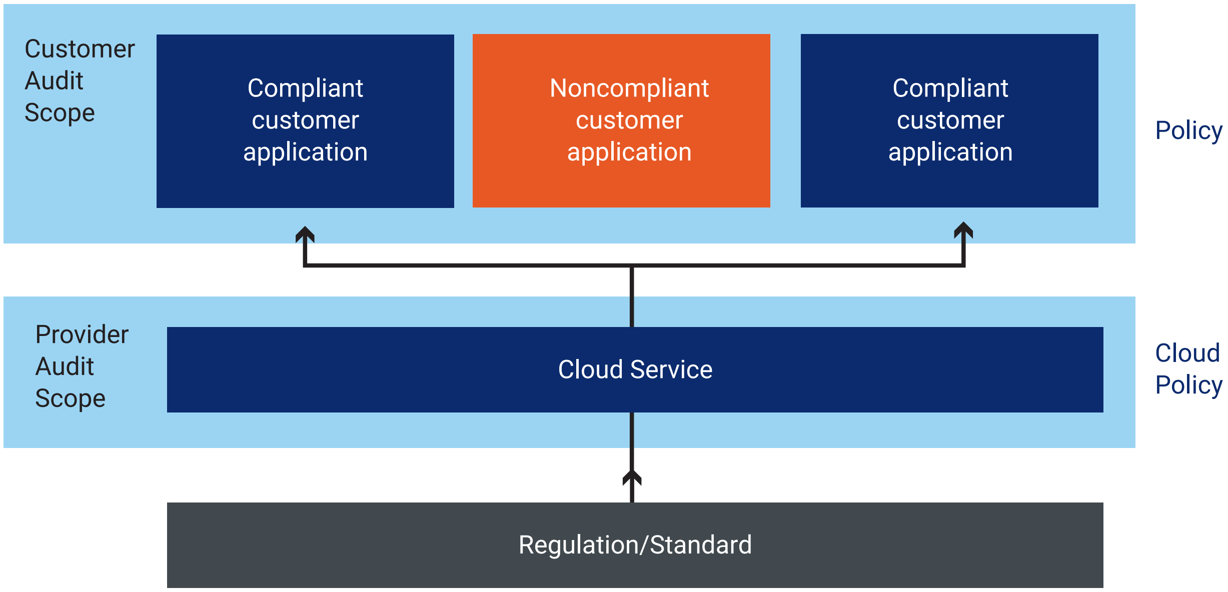
Customer Audit Scope
Compliant customer application
Noncompliant customer application
Compliant customer application
Policy
Provider Audit Scope
Cloud Service
Cloud Policy
Regulation/Standard

Source: Cloud Security Alliance, "Security Guidance for Critical Areas of Focus in Cloud Computing v4.0," 2017, https://downloads.cloudsecurityalliance.org/assets/research/security-guidance/security-guidance-v4-FINAL.pdf

5.7 Understanding the Audit Context

The initial stage is a foundational part of an audit. It is about establishing the organization's context, obtaining a full understanding of the organization, its interested parties and its environment—i.e., how it operates, and its involvement with events, conditions and types of transactions that might impact the cloud environment. Ultimately, the objective is to confirm or validate an understanding of the organization's current state, and how the cloud audit plan fits into the existing auditing approach.

Note: It cannot be stated often enough: The auditor will be assessing the organization effectiveness of performance—not whether it is purely right or wrong—and whether the cloud policy, such as it is, supports the business risk appetite.

It may help to think of this initial stage of audit as laying the foundation for a house. The organization should have a sense—beyond the auditor's assessment—of the most critical cloud applications it is consuming. For example, if it is running a significant part of its operations in the cloud, the risk is not limited to cybersecurity alone; the cloud may also represent a risk to customer service, or to sales. This exercise should highlight the criticality of the cloud service, along with dependencies and risk.

Developing an initial data security process flow may be helpful at this stage to facilitate the analysis of the data flow, and who handles the data, along with details behind the inputs and outputs. In addition, tools such as heat maps can be helpful to indicate the importance of a service and the consequences if it were unavailable. As part of understanding the operations, it can be useful to interview top management to understand the organization's overall strategic direction. For example, if the business anticipates it will triple sales in the next quarter, then its infrastructure might be affected, which could bring capacity management into scope.

5.7.1 How Cloud Audit Planning Fits Into the Existing Company Approach

The organization may have a multifaceted approach to audit preparation, depending on differing requirements to demonstrate compliance with particular regulations or standards. The controls it applies may be different for its financial systems, for example, than for a card processing instance. The cloud audit plan must be integrated into the existing company approach that evaluates the individual regulatory or standards requirements.

During this planning stage, the goal should be to have a full blueprint of the organization's cloud activity, informed by an understanding of what the organization does. When building a cloud audit preparation plan, it's important to build an inventory of the current state of cloud activities in the organization and how they relate or integrate with its other operations. This covers both organically grown cloud activity—that is, where internal IT exercises centralized command-and-control over the cloud, and where it is running policy and engagement with CSPs—and renegade activity, sometimes known as shadow IT, as the organization may be dealing with other CSPs and might not be aware of it. This stage should also encompass an understanding of whether the chosen CSPs are further outsourcing the cloud activity (which could represent third-, fourth- or fifth-party risk).

While carrying out the inventory, the organization needs to be mindful of data location geographically—where the CSP servers are physically based, and what national jurisdictional regulations govern those servers. (For instance, the EU General Data Protection Regulation covers not only entities operating in the European Union, but also entities processing data about EU citizens, regardless of where they operate.)

This building-and-planning activity should include some inventory of the organization's service level agreements (SLAs), together with any policies relating to cloud activity. SLAs are one of the most important sources of evidence when assessing risk. This exercise establishes an understanding of the organization's shared responsibilities with respect to the cloud services and CSPs it has engaged. It informs the risk analysis and establishes shared governance responsibilities, which is a key element with cloud services.

To prepare for a cloud audit, the organization needs to know how many SLAs it has, with how many CSPs. In addition, it should know what specific service levels it has contracted for and what is involved in each SLA. What rights does the customer have for auditing, and what reporting and attestations does the CSP provide? For example, in the US, a SOC 2, ROC AOC for PCI or ISO/IEC 27001 would confirm the system is SLA-driven.

Even if the cloud is entering its maturity phase, many organizations may not yet have specific cloud-related policies. This is a sizable gap, and one of the critical deliverables of an auditor is to identify gaps in policy—even if a certain control is inherently understood. The auditor may recommend prioritizing gaps that could be materially important; for example, if the organization needs a written policy concerning shadow cloud activities.

From a cloud perspective, what are the unique regulatory requirements for cloud computing that apply to the organization?

Tip: When carrying out an inventory or current state analysis, some important elements to capture include listing what current audit or attestation reports the organization is required to have. This is an opportunity to align those reports with internal policies and SLAs with providers to see where there may be further gaps. The auditor's report should state those recommendations.

Understanding the context for a cloud audit also covers identifying the people in the organization who know the environment. This exercise should also consider the current maturity level of cloud competencies in the IT organization.

By reviewing past reports at the stage of assessing risk factors, the auditor will get an idea of where some of the problems may be, which helps in prioritizing the risk. If the reports are recent, the auditor may be able to rely on the evidence they contain or identify what additional information to gather for carrying out the risk assessment.

One point of difference: By virtue of the speed at which it is developing, the cloud sector is moving toward a continuous auditing model, so the organization may take an approach to cloud auditing that differs from its existing audit plan by adopting continuous auditing.

5.8 Audit Building/Planning

Conceptually, a cloud audit is not much different from a traditional IT audit. Many fundamentals of an assessment or audit still apply, just as they would in a traditional IT environment. The auditor or assessor is viewing what the company has done through the lens of the controls it has put in place in the form of a policy and procedural review. (See **figure 5.8**.)

Figure 5.8—Audit Process

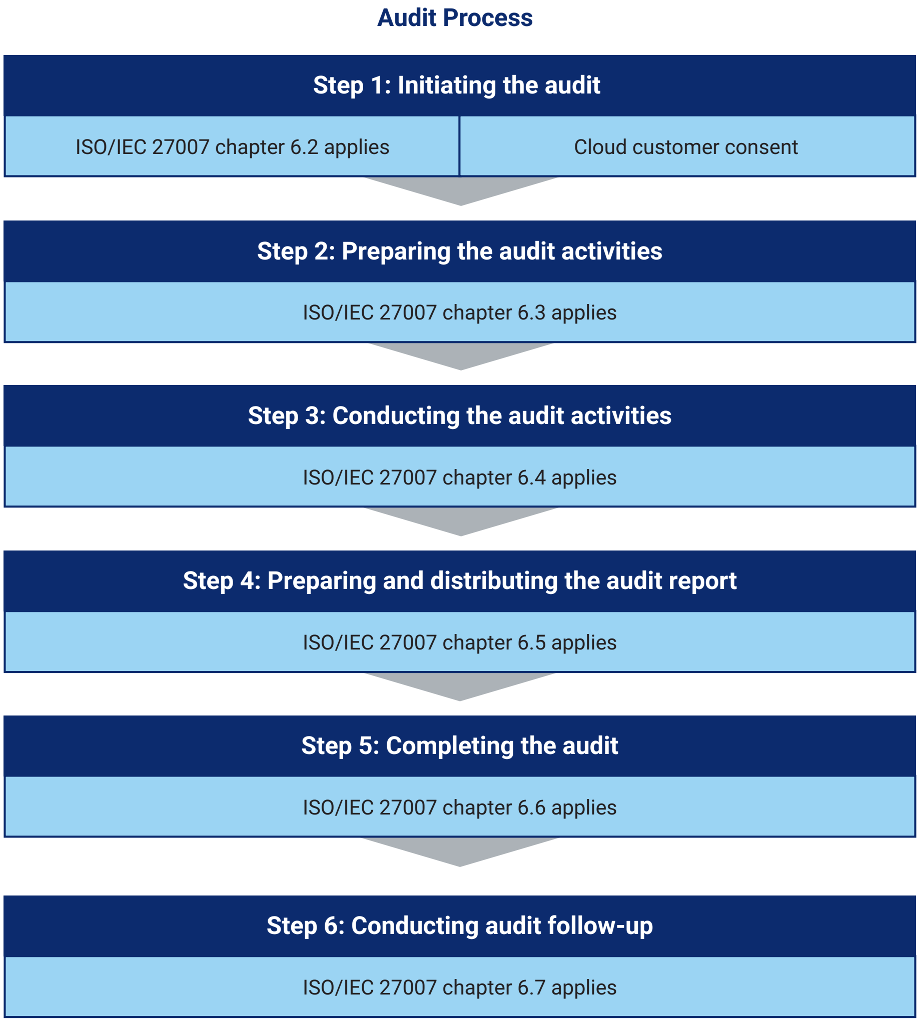
Audit Process
Step 1: Initiating the audit
ISO/IEC 27007 chapter 6.2 applies
Cloud customer consent
Step 2: Preparing the audit activities
ISO/IEC 27007 chapter 6.3 applies
Step 3: Conducting the audit activities
ISO/IEC 27007 chapter 6.4 applies
Step 4: Preparing and distributing the audit report
ISO/IEC 27007 chapter 6.5 applies
Step 5: Completing the audit
ISO/IEC 27007 chapter 6.6 applies
Step 6: Conducting audit follow-up
ISO/IEC 27007 chapter 6.7 applies

With traditional IT audits, change management, logical access and computer operation controls are part of the scope. Where the cloud introduces differences is in placing a greater focus on third-party dependencies, and therefore a greater focus on boundary definition. It is a more complicated view of what is traditionally considered to be a boundary, and it implies security everywhere, as opposed to security in one location. Although some dependencies are relatively common and well understood—such as Electronic Data Interchange (EDI) integrations or dependencies on third-party infrastructure providers, or on ISPs—some boundary definitions are less well-defined, and security can be impacted by something that is two or three steps removed from the organization. With the cloud, the adage trust but verify applies.

5.8.1 Scope of the Assessment/Audit

When it comes to building and planning an assessment or audit, defining the audit objective and conducting scoping of the audit engagement are by far the most important exercises to complete. Organizations and their auditors must have a sense of its environment, its objectives, and the purpose its external stakeholders are trying to achieve, the regulatory environment in which the organization operates, and which controls are necessary to achieve compliance.

A discussion on scope should begin with the audit objective the organization is trying to accomplish, using that objective (for instance, receive PCI DSS certification to handle payments data for an online store) as a primary input into defining the scope of its auditable environment. This scope should be tailored in relation to the organization's business needs, the structure of the organization, its location, its information assets and technologies, and any overlapping audit scopes (for instance, ISAE 3402/SSAE18 scopes that cover different environments).

When conducted well, the scoping phase ensures no stone is left unturned by identifying the inventory the auditor needs to focus on, whom to talk to in the organization, and what evidence to gather.

In order to get to an effective engagement scope, the scoping discussion should review the audit objective, the components that make up the scope (context),[242] the previously mentioned tailoring, and the entities and dependencies that form the scope.

During the scoping phase, it is important that the internal or external auditor understand the expectations of the sponsors commissioning the audit. The auditor should set the expectation that the exploratory activity will identify immediate dependencies, which may lead to a recommendation for a broadening of scope. This can be a difficult conversation and may influence the scope of the audit. Scoping discussions about audits can be contentious. There needs to be clarity between the stakeholder and the auditor from the outset, with agreement on the focus of the audit—e.g., a software development assurance process, the SaaS platform it uses, its security operations—and the items and processes that are in and out of scope. This allows for a clear road map and avoids scope creep.[243]

The auditor should identify if there are multiple stakeholders who will receive the same copy of the report, or if there will be different summary statements, requiring different levels of detail, e.g., will the audit report contribute to an external certification the organization wants to obtain?

At this point, it may be useful to build a RACI chart or responsibility matrix showing what controls the user organization is responsible for and what the CSP will cover, to get a better picture of scope and shared responsibility.

The big picture context can be summed up as follows:

- What is the organization's use of cloud services?
- Where is the organization doing it?
- Internal issues: factors under the direct control of the organization
- External issues: factors an organization has no control over, but that it can anticipate and adapt to
- Who carries out this activity for the organization?
- What policies support this?

Data flow diagrams and supporting contextual information can help to build an understanding of where the boundaries lie within the organization, the connection to service level agreements, the delineation of shared responsibilities, and how data flow between applications and CSPs. This establishes what is specific or unique to the cloud that the organization needs to consider.

When determining what is in scope, it can help to think of a series of borders defining boundaries (**figure 5.9**). The smallest border represents the internal controls and policies over which the organization has oversight and control. A

[242] An organization must consider both the internal and external issues that can impact its strategic objectives and the planning of the security system. It is also a means to detect risks and opportunities regarding the business context.

[243] Scope creep occurs when the audit or project stretches beyond its originally defined scope, e.g., when the scope of a project is not properly defined, documented or controlled.

larger border represents a CSP, and any further borders represent additional third-party interactions. Some controls outside the first boundary may be defined in the SLA.

Figure 5.9—Scope Example

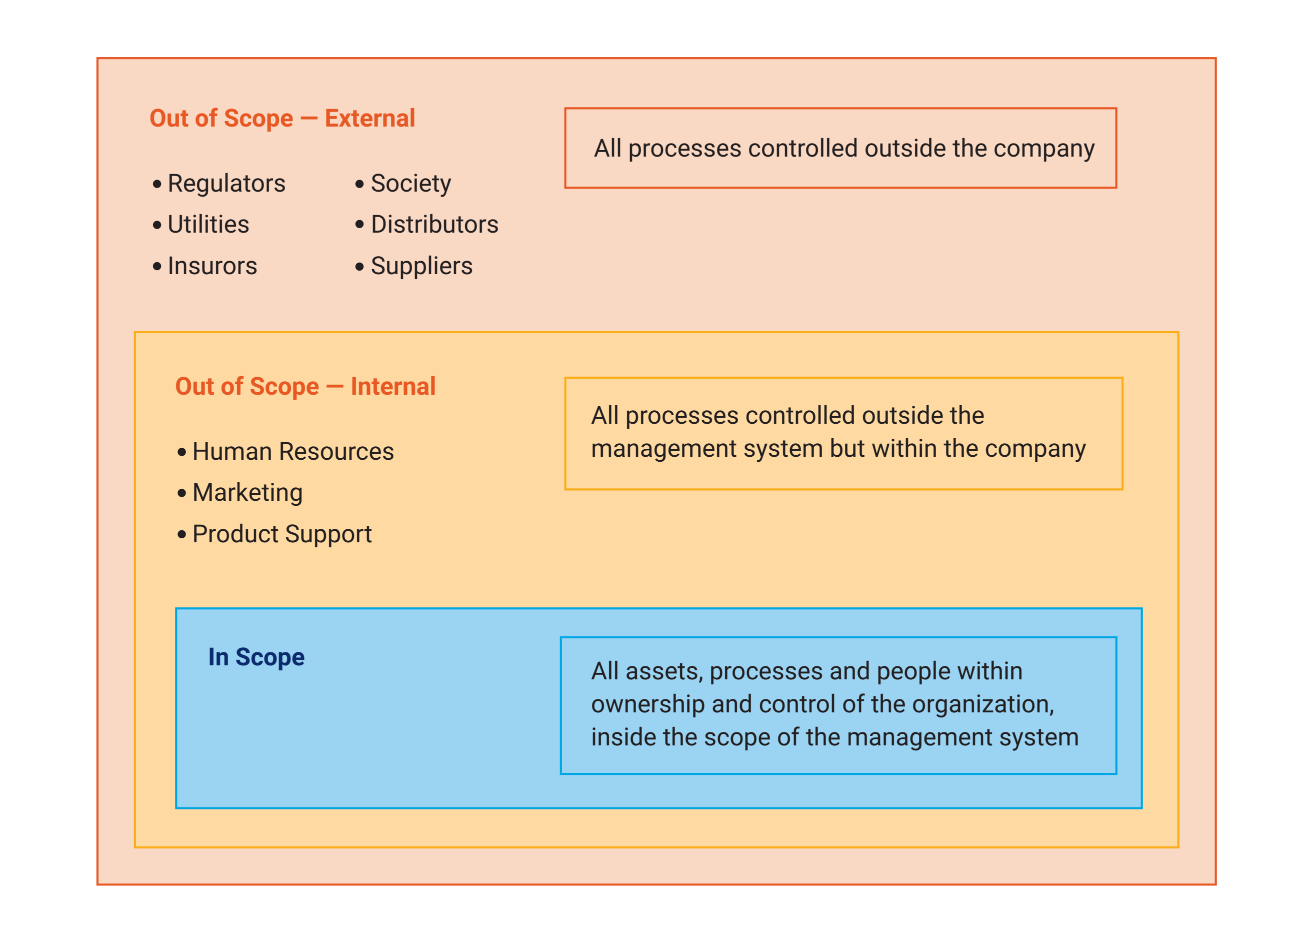
Out of Scope — External
• Regulators
• Utilities
• Insurors
• Society
• Distributors
• Suppliers
All processes controlled outside the company
Out of Scope — Internal
• Human Resources
• Marketing
• Product Support
All processes controlled outside the management system but within the company
In Scope
All assets, processes and people within ownership and control of the organization, inside the scope of the management system

When considering the first boundary, organizations should focus on what services they provide and what core web components are involved. It can be useful to think in terms of the user interface (UI) or user experience (UX) in the case of a command-line web application. In the classic three-tier web architecture, where is the data that are presented to the customer? Behind the UI or UX is the database or storage tier. The first tier or boundary does not cover third-party interactions. Then, the questions to ask include these:

- What external systems can impact the security of the boundary?
- What systems external to the boundary hold data that the customer uploads along with any sort of data that are directly extrapolated? (For example, a business continuity plan site or a failover site should be in scope if the organization will store customer data there, because it automatically would become a key piece of the system if the primary system should fail.)

Other questions to consider in the scope:

- Where else does customer data reside, besides on internal systems?
- What is the organization running that could impact the security of the system?
- Many CSPs and a lot of organizations have corporate-level systems, such as a central SOC or SIEM, that manage security (e.g., access control or centralized logs). They may have security-related data feeds coming into one place so they can monitor overall hygiene. Are these shared facilities accounted for?

First, gather this information:

- What service the organization sells/what the customer gets
- Any shared services or access controls that also need to be in scope, even though no customer data are shared

Another layer to consider including in the scope may be external to the customer system, such as security tooling or a white hat application to run automated security testing. The organization may have external tie-ins that impact the security of its environment. If building on one of the large infrastructure providers, the organization may believe that its CSP is handling a particular control. However, under the shared responsibility model, the organization consuming the service must still configure logins, single sign-ons, etc.

Some CSPs, particularly the more mature and developed ones, have identified common misconfigurations and released ways for customers to prevent data from mistakenly being publicly accessible. Some of them provide features to carry out configuration checks and methods to prove the landscape of the customer's environment is secure. The auditor should provide interpretation on what those controls are, if they exist, and who is responsible for implementing and maintaining them. In short, the auditor should look for evidence that the organization has carried out its own due diligence around its obligations, particularly those subject to regulations.

Once a cloud audit scope expands beyond a certain level, it effectively becomes unsustainable, mainly because if no additional resources are added, the time frame for completion of the audit is jeopardized, and there is no longer alignment between the audit objectives stated during planning and the work being done. The organization needs to draw a line between what is realistically achievable and what is not to avoid scope creep. If the organization must abide by a regulation that requires a password policy, the CSP should support the customer's ability to maintain the same type of password configuration.

Note: From an auditor's perspective, the preceding point is not a critical finding, but a possible opportunity for improvement.

While key controls that are within the context of the organization are considered in scope (unless there is clear justification for why some are not applicable), some controls that are outsourced, although technically not directly applicable, still fall under supplier relationships, and an organization will have to show how it performed due diligence to ensure that the third party is meeting its obligations. An organization may decide that if it requires CSA STAR and it chooses to buy a service from a particular CSP because it has attained CSA STAR Certification, this may suffice as proof of due diligence.

5.8.2 Shared Responsibility Model and the Role of the SLAs

One of the least understood but most impactful changes to cloud risk management and governance is the concept of the shared responsibility model (**figure 5.10**). While individual CSPs may employ similar definitions and implementation choices, the techniques and best practices each CSP uses may vary widely for establishing, operating, modifying, monitoring, and reporting on controls and configurations.

Awareness alone does not solve the accountability and responsibility requirements. Multicloud governance requires mastery of each CSP configuration and setting. Validating the accuracy and integrity of CSP configurations and developing relevant software and tools capable of monitoring their integrity and performance are major short-term challenges for cloud customers.

With the cloud, organizations do not have direct access to, or the ability to test, a CSP physical infrastructure. However, they can use data provided through SLAs. Both the customer and the CSP will refer to such data in a service event such as an outage or a degradation of service. Cloud SLAs are important to several stakeholders:

- For technology operations, the SLA defines what they can expect from the cloud provider's service. The technology operations group will typically manage any deviation from the service.
- For the sourcing and finance functions, the SLA is linked directly to financial spend and sourcing due diligence.
- From the governance or control functions, the SLA is considered the bible for interpreting events, particularly for assessors and auditors.

When examining a CSP SLA, an auditor is determining whether the agreement excludes any operational matters that are material to the operation of the service.

5.8.3 Analyze the Data Flow and Overall Architecture

Figure 5.10 shows the shared responsibility model.

Figure 5.10—Shared Responsibility Model

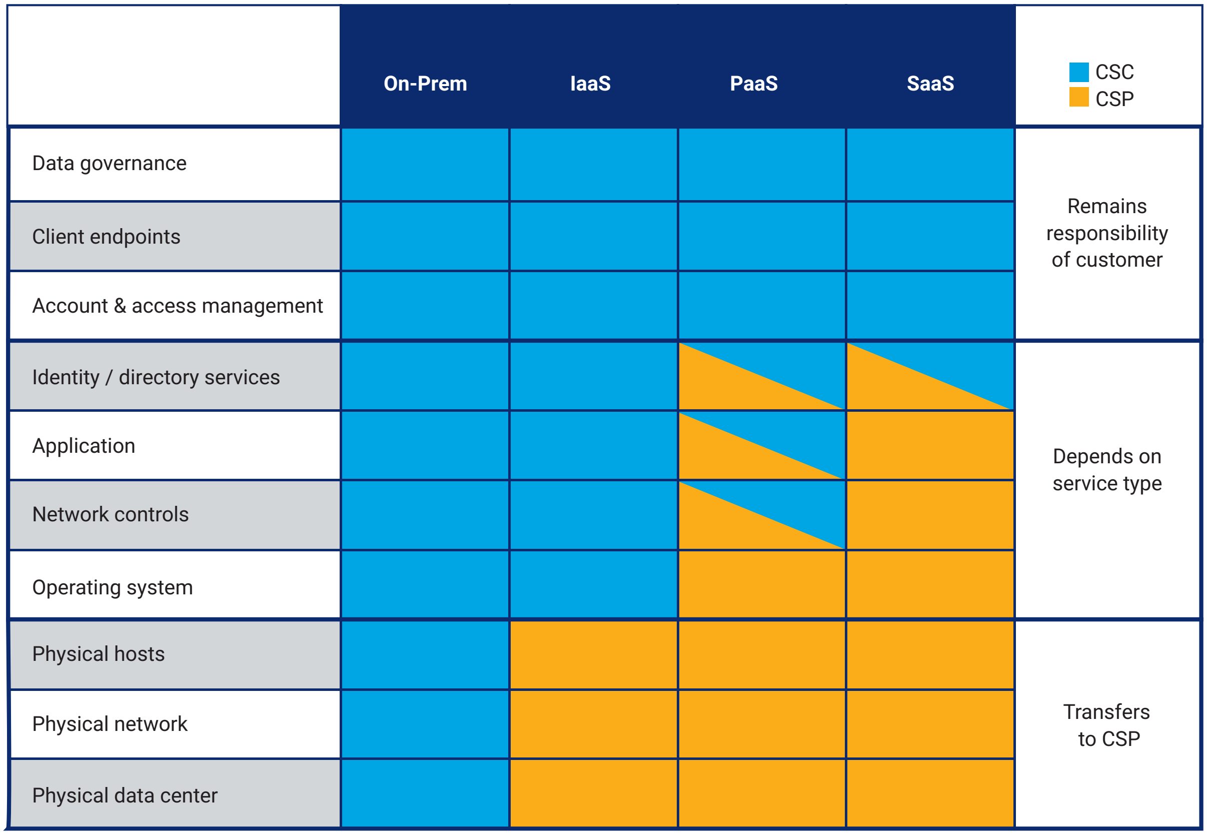

	On-Prem	IaaS	PaaS	SaaS	🟦 CSC 🟧 CSP
Data governance	🟦	🟦	🟦	🟦	Remains responsibility of customer
Client endpoints	🟦	🟦	🟦	🟦	
Account & access management	🟦	🟦	🟦	🟦	
Identity / directory services	🟦	🟦	🟦🟧	🟦🟧	Depends on service type
Application	🟦	🟦	🟦🟧	🟧	
Network controls	🟦	🟦	🟦🟧	🟧	
Operating system	🟦	🟦	🟧	🟧	
Physical hosts	🟦	🟧	🟧	🟧	Transfers to CSP
Physical network	🟦	🟧	🟧	🟧	
Physical data center	🟦	🟧	🟧	🟧	

This segment is intertwined with the scoping stage, as the scope is informed by the data flow, which itself is informed by the customer experience. When discussing data flow with the customer organization, a useful perspective for the auditor is to examine what the organization is selling, where the customer data sit, and what internal data sets impact the security of the environment.

It can be useful to arrange a user-experience (UX) meeting and ask for a standard demo of the organization service, focusing on security discussions. The auditor asks the relevant individual or team in the organization to walk through the system: what it provides the customer and how the customer interacts with it. This helps the auditor to gain a sense of how a customer interacts with the organization's system by uploading, downloading, transmitting or processing data on it. After gaining an understanding of how the customer interacts, the auditor can look at the scope through an auditing lens: for example, at the back end of the system, if the customer uploads a document, where does that land? The auditor can trace all the ways the user interacts.

The auditor can get a sense of the customer-facing architecture (and how it informs the boundaries of responsibility), and take it a step further by thinking of the control context. The auditor can get a picture of the security system and understand how the enterprise has defined the boundaries of the system when talking about the cloud to get a picture of where a CSP might rely on a particular standard for attestation.

While the organization will originally stipulate what it wants tested as part of the audit or assessment, it is still the auditor's responsibility to trace the data and track the security architecture, which is in line with the original agreement of what is in scope, per the process defined in section 5.8.1.

5.8.4 The Assessment/Audit Criteria

Different regulations have different requirements and set different auditor expectations, including the depth of testing. They possibly also extend to the auditor's independence from the enterprise. For example, PCI-DSS places more focus on cardholder data than ISO/IEC 27001, which is concerned with strong risk management, while FedRAMP, the US government standard for cloud services, is more stringent than CSA STAR. The controls may not be the same, therefore, depending on which criteria need to be applied.

In addition to becoming familiar with the standards s, the auditor should consider two things:

- What does the standard in question require from the auditor?
- How strict should the auditor be when evaluating controls? (For instance, some frameworks accept risk-based justifications or mitigating factors when an auditor evaluates a control, while others require the auditor to use a binary implemented vs. not in place approach.)

Cloud adoption is a journey. If an organization is starting out on its adoption of cloud and is relatively immature, then it may be mapping controls from an established standard or regulation in order to demonstrate compliance.

5.8.5 Roles and Responsibilities of Individuals

This is a list of whom the assessor or auditor needs to work with in the organization as part of carrying out the audit or assessment. It is part of the scoping exercise, and the list can then become the interview schedule for the audit itself. Questions to consider include, but are not limited to the following:

- Who has logical and physical access to the system?
- Who should not have access?
- Who can interact with the customer data?
- Who has control of security?
- Who is responsible for interacting with third parties (e.g., API calls)?
- Who performs background checks?
- Who ensures that the appropriate people set policies?
- What are the controls that deal with people in the organization, e.g., security awareness and training?

5.8.6 Competence of Individual(s)

The auditor must be familiar with cloud concepts, virtualization, abstraction, software, infrastructure and Platform as a Service, and how APIs work.[244] Then, depending on the regulatory requirements the organization must meet, the auditor may need to have specific knowledge or qualifications, such as QSA if carrying out PCI audits, or a CPA for carrying out SOC 2 attestations or competencies in line with ISO/IEC 27006. The auditor should be up to date with certifications against the relevant standards, and possibly also have undertaken some CSA training, such as CCSK, to prove competence in the cloud and facilitate how one might apply assessment and auditing specifically for cloud services.

5.9 Audit Execution

The key to executing an audit is to engage in cross-functional, business-wide communications that ensure transparency at every point about the timing, planning, preparation and outcome. This approach guarantees that all stakeholders are aware the audit is going on, and are informed of its purpose, and the expectation and objective from the audit both during and after completion. It is essential to have everything documented, with questions and follow-ups explicitly explained. There should be closing communications at the end of the audit process explaining next steps and thanking every participant involved for providing input.

5.9.1 Opening Meeting/Kick-Off

The auditor should send out invitations before the initiation meeting to invite and copy the selected participants, making sure to include the relevant audience and the audit stakeholders, together with others as senior in the organization as possible. When presenting, the auditor should be clear about the reasons why the participants should be there, helping them to understand the value of their participation. It can be difficult to put a monetary value on compliance alone, so it may be necessary to make a case expressed in terms of return on investment, annualized risk reduction, or how the audit impacts the organization's ability to do business. For internal audit over IT controls on cloud activity, the case might be about following internal corporate policy.

The kick-off meeting should be concise and inclusive—it generally should not exceed an hour. It may be helpful to think of it as an executive summary. Stoplight charts can quickly draw stakeholders' attention to key areas of focus for the audit itself.

The objectives of the kick-off meeting:

- Outline the plan at a high level the plan and get validation for that plan.
- Introduce key players within the auditee organization and note available escalation paths.
- Clearly identify the scope of the audit.
- Identify additional stakeholders/people that need to be interviewed.
- Identify any risks and dependencies from people out of the office to dependent service providers like data centers or other cloud service providers.

Certain audit frameworks, like ISO/IEC27001, dictate topics that need to be covered in the kickoff meeting.

[244] When an audit team is carrying out the audit/assessment, it must have the necessary credentials as a unit. This means each member of the team does not necessarily need to have all the same credentials.

5.9.2 Communication Before/During Audit/Assessment

Having established initial communication before the first meeting, the auditor should continue the process by publishing the plan at that kick-off meeting. The auditor will communicate relevant parts of the plan as the audit proceeds, reconfirming the plan all along, highlighting an aspect and drawing attention to the names of individuals involved at that particular stage. A key element is letting the auditees know and giving them time to prepare. An audit is not a surprise or an ambush. Auditees appreciate having the chance to prepare. An auditor who knows what evidentiary documents are needed should give the auditee the opportunity to put them together in advance, saving time and effort for everyone. Having those documents in advance gives the auditor the advantage of arriving at the first meeting with deeper and more thoughtful questions.

5.9.3 Sampling

Audit sampling by definition is the application of a procedure against something less than 100 percent of the population of an element, such as a transaction or balance in financial audit,[245] but more appropriately against an attribute (control) or subset of systems within a cloud service.

Sampling will depend on a combination of requirements regarding regulations, internal and external standards, policies and scope. However, the auditor's judgment is equally if not more critical. If there is no data population big enough to carry out statistical models, then the auditor's judgment is more valuable. Many audit firms take what is typically high-level scoping guidance from audit standards and apply their own additional guidance so that the audit teams test in a consistent manner. A key consideration is what is defined as sampling risk—the risk that the sample selection is such that the auditor's conclusion would be different if sampling were not applied. There is sampling risk[2] any time a sample approach is utilized, and there is more than one approach to sampling.

Nonstatistical or Judgment-Based

Since priorities vary from one organization to the next, and the risks will be different for each, the auditor should use nonstatistical-based sampling that takes into account the cloud services that are most important to the organization such as business-critical systems. It is good practice to check the controls around the perimeter and those guarding the most important systems and data in the organization, but the auditor should assess the organization on its own merits. Services most likely to be affected by a security incident could include, for example, web servers prone to defacement. Further, email servers, domain controllers and firewalls are regularly exploited. Database servers are vulnerable to attack. This is not an exhaustive list.

Statistical

There is a misconception that statistical sampling does not require judgment. That is not the case and it is the reason standards bodies such as the AICPA denote that professional judgment is required for all sampling approaches. For example, statistical sampling is often performed in the following scenarios:

- A sample of virtual machines, storage buckets or other elements that may be voluminous in the cloud (For example, an auditor selects a sample of 25 system components across the application, web, database, etc.)
- A sample of 25 application changes (or change tickets)
- A sample of 25 new hires for onboarding controls
- A sample of 25 terminated employees for offboarding controls

[245] AICPA, "AU Section 350 Audit Sampling," www.aicpa.org/Research/Standards/AuditAttest/DownloadableDocuments/AU-00350.pdf

The sample of 25, which comes from a commonly practiced approach of taking the lesser of 25% of control instances or 25, is just that—a practice. A prime example of judgment coming into play is the use of centralized controls, such as configuration management tools, that would allow the auditor to reduce the sample size or what was reviewed. For instance, an auditor who had confidence that configuration tools such as Puppet were used could focus on the Puppet configurations and review the sampled systems for the presence and operating effectiveness of the Puppet agents, as opposed to testing passwords and other parameters that the agents may enforce.

5.9.4 Fieldwork, Data Collection and Documentation Review

Data-Evidence Collection

All audits, by definition, require the collection, management, use and storage of audit evidence. For successful evidence collection, an audit involves a mix of techniques used interchangeably: visual observation, examination of records, walk-throughs, reperformance of controls and stakeholder interviews. They all contribute to a collective sum of all the parts, which ultimately provides the auditor with a solid basis for a conclusion regarding compliance.

The auditor should begin by reviewing the documentation printouts and any other relevant data that was created when implementing the security system. The audit scope should align with the scope and context (inputs and outputs) of the organization and verify what is being audited. The auditor should observe how the system works in practice by speaking with front-line staff members. The auditor should perform audit tests to validate evidence as it is gathered and review internal data and reports from the organization to validate integrity of the data. The auditor should sort and review the evidence collected in the audit in relation to the organization risk treatment plan and control objectives. Occasionally, this analysis may reveal gaps in the evidence or indicate the need for more audit tests.

With respect to evidence collection, ISO 19011:2018 states that “information relevant to the audit objectives, scope and criteria, including information relating to interfaces between functions, activities and processes should be collected by means of appropriate sampling and should be verified, as far as practicable.”

Only information that can be subject to some degree of verification should be accepted as audit evidence. Where the degree of verification is low (for example, information collected from personnel interviews without a corroborating account from another person or an accompanying artifact) the auditor should use professional judgment to determine the degree of reliance that can be placed on it as evidence. Audit evidence leading to audit findings should be recorded. If, during the collection of objective evidence, the audit team becomes aware of any new or changed circumstances, risks or opportunities, it should address them accordingly.

In addition, the guidance, specifically for SOC 2 (similar to ISO/IEC 27001), is to ensure that there is reasonable assurance gained by the CPA firm in order to issue an opinion, whether in the form of Type 1 or Type 2. The evidence supporting the opinion must ensure that conclusions can be made whether controls are in place (to meet Type 1) or are in place and operating effectively (to meet Type 2). Peer reviews are performed by independent assessors that sample files and ensure that processes are performed accordingly. No specific evidence requirements are published, but if the audit was done poorly or improperly the repercussion would be losing a CPA license.

It is prudent as a best practice to follow ISO 19011 because that is typically the “standard of care” that would be accepted in an authoritative document as a baseline. Chain of custody cannot be analyzed in stages—the entire process must be taken into consideration. Chain of custody in legal terms is the chronological documentation or paper trail that records the sequence of custody, control, transfer, analysis and disposition of physical or electronic evidence.[246, 247] As stated in ISO/IEC 19011, the evidence should be verifiable, and for this reason (as well as for IT forensics and admissibility in a court of law), a chain of custody shall be maintained.

[246] EDRM, “EDRM Model,” www.edrm.net/edrm-model/

[247] Legal Dictionary, “Chain of Custody,” 9 March 2019, https://legaldictionary.net/chain-of-custody/

Due to the automation of most processes and controls associated with cloud platforms and technologies, it can be challenging to ensure that the integrity of the audit trail over data identification, data analysis and testing is maintained throughout the audit process. This is because ownership of the data transitions between the CSP and the customer conducting the audit. This is especially important when the customer is using the cloud provider's tools and processes to collect, process, store and manage the audit evidence.

One classical audit tool used to ensure a well-documented audit trail is the turtle diagram (**figure 5.11**). A turtle diagram is a process auditing and improvement tool (schematic visual representation) of the key elements that make up a process, including the resources needed to achieve its purpose. Within the anatomy of a turtle diagram is the visual display of key process characteristics such as inputs, outputs (expectations), criteria (metrics) and other high-level information used to assist relevant auditors and process stakeholders in the effective execution of an audit, and to facilitate improvement of the process itself.

Figure 5.11—Turtle Diagram

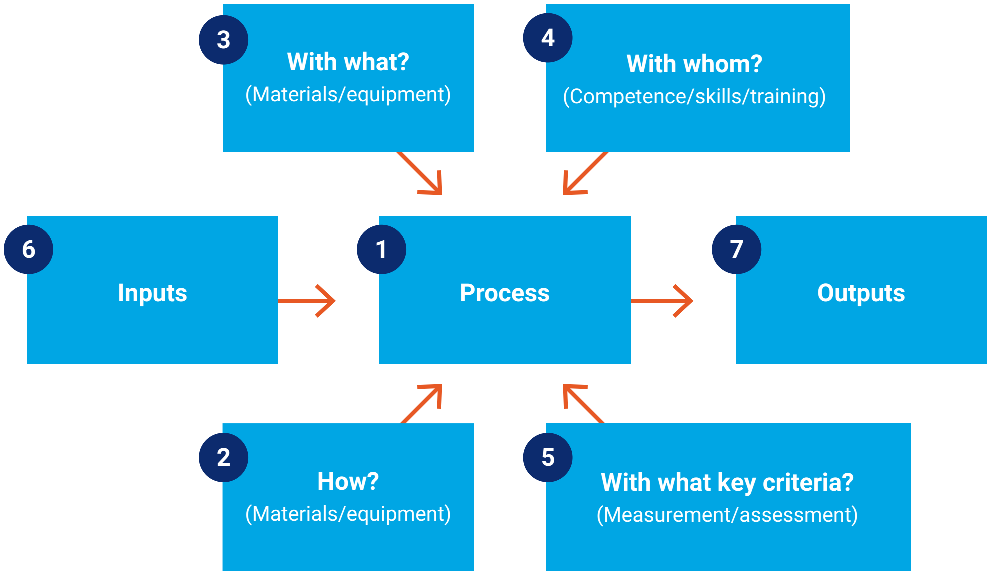

Section	Details
1	Enter process or support process name.
2	Enter details of linked process controls, support process, procedures, methods, etc.
3	Enter details of the technology (including test equipment), computer systems, software used in the process.
4	Enter resource requirements, paying particular attention to required skills and competence criteria.
5	Enter the measures of process effectiveness, i.e., matrix and target.
6	Enter details of the actual input. This may be a document, materials, tools, schedule, etc.
7	Enter details of the actual output. This may be a product or document, and it should be linked to actual measure of effectiveness.

NIST SP 500-291 identifies available information on current standards, standards gaps and standardization priorities. While not the only resource, it is an excellent source that walks an auditor through the plethora of auditing challenges with multiple cloud services and provides guidance including case studies and use cases.

5.9.5 Generating Audit Findings

The auditor or assessor should first generate draft findings. This is to document observations, which will be based on context, in addition to presenting fact-based information. The auditor should draft findings in a way that is impersonal, possibly stating them in the form of questions be validated with the stakeholders.

Tip: A note on wording: Passive voice helps when writing the audit report. The aim of the report is to identify conditions and states, not people. For example, an auditor who finds that not all machines are updated with the same patches might present this draft finding as follows: "Confirming I find the Windows 10 machines patched to the current level, but not all Windows 7 machines are patched to the same level." The auditor may present this finding in another way: "During the audit, it was confirmed that the Windows 10 machines are patched to the current level, but not all Windows 7 machines are patched to the same level."

Note that draft findings are discussed fully with the auditees and may be mitigated, i.e., the auditor may discover mitigating factors or compensating controls when presenting draft findings.

At this stage, anything that might cause the organization to fail the audit is couched as a potential finding. This part of the process is important, because the auditor does not want to issue a report that is contested. Stakeholder buy-in is essential to getting agreement that a particular situation exists. The auditee must confirm any stating of fact, even if there may be disagreement over why it is in the report, or the extent of its impact.

The audit team should ensure there is sufficient, appropriate audit evidence to support the results reported, thereby ensuring that everything reportable is reported, and everything reported is potentially reportable by a third party or certified body as well.

Information gathered during the audit may be sensitive and confidential. Therefore, all computers used during this project must be equipped with file-level encryption programs. At the end of the project, the audit team should delete all files and information associated with the review.

5.9.6 Continuous Auditing

See section 6.8.1 to review continuous auditing and continuous monitoring. With IT service management platforms, continuous auditing works most effectively as self-service, so that as part of the organization's process it generates evidentiary documents (reports, etc.) continually on a daily or weekly basis to enable auditing at any point in time. This requires audit tools and policy management tools that validate compliance and policies, or gather information for the audits. Continuous audit tools should be as automated as possible, collecting compliance evidence through logs, event databases, cloud API queries, configuration analyses, and testing. Topics such as backup and recovery, vulnerability assessments, static and dynamic code analysis, logging and filtering, cryptography and authentication are well-suited for this approach. On the other hand, procedures, governance and documentation quality are much harder to evaluate through automated continuous auditing.

Note: High levels of fully automated audit and compliance are found in highly mature organizations that can identify vulnerabilities and threats and manage remediation without human intervention.

Continuous auditing is not beyond an organization, by any means. However, it requires a lot of work, and generating continuous reports needs to be part of standard operating procedures. It is procedural; while it can be automated, most organizations today will be leveraging a combination of tools and human resources to inspect controls. From an internal perspective, continuous auditing would have to be through automation and through tools communicating with each other via APIs, security information and event management (SIEM), or a central dashboard. The governance risk and compliance tool may prompt users to generate evidentiary documents and scan for vulnerabilities.

5.9.7 Reporting and Report Distribution

The purpose of an audit report is to inform decision makers at one or more stakeholders (a client, a regulator, a health and safety body, a standards body, etc.). Internal audit reports can be requested by an organization's board of directors. Generally, the internal audit function has a program of work with a high-level audit plan that reflects regulatory requirements or principles.

Report Types/Sections

Most audit reports follow a standard format, detailing the following:

- The scope (what was audited—and maybe a rationale for what was and was not in scope)
- The subject matter experts consulted
- When and how the audit was conducted
- Extracts from relevant artifacts and associated reports
- Observations and recommendations

Typically, the audit report is a mix of quantitative and qualitative commentary by the auditor.

The quantitative aspect expresses levels of compliance; the qualitative element will share the nature of the finding or the observation. For example, the report may state that the audited entity is performing tasks in a way that meets the requirements of a standard, but its software tools require manual checking.

The report normally presents some level of evidence that flows from the stated methodology to assert a fact. An auditor who was taking a sampling approach may present verification of three out of 100 transactions, for example.

The report will include findings, observations, or recommendations for action. The auditor may attach monitoring or management guidance, which can include target closure dates for corrective actions to reflect finding sensitivity, and suggested owners to ensure the appropriate person is identified to resolve the recommendation. It may also suggest other areas that may be audited in the future, e.g., scheduling a future audit when a department identifies plans to implement a certain control in a certain time period.

Many internal audit reports have an area for the auditee to respond—either to accept the audit point, outline what corrective action it will take and when. If the auditee disputes or rejects a finding, that will involve a series of meetings until a resolution is found.

Tip: It is not uncommon in the audit practice to review audit findings with the groups being audited before the report is finalized. Then, if there is a dispute, resolution can be found prior to finalizing the report.

Sharing of Audit Results and Artifacts

During the audit work, artifacts including working papers and ancillary documents—such as policies, or logs the auditee presented as evidence of performing a certain control—are often discovered. Statements made by management in response to specific inquiries (or through financial statements) are referred to as management representations. If verbal, these representations should be captured in writing. Since auditors often must rely on these statements, it is common practice in entrance meetings, and possibly in an audit engagement letter, to include wording reflecting these representations. Management representations are not a substitute for audit work. They represent management's view and are subject to challenge as appropriate. Any history of discrepancy between management representations and ground truth can be reiterated in the form of observations and investigated during fieldwork.

The auditor should be sure to capture statements the organization made, or meeting minutes, and to organize any files or links. Sharing minutes of a meeting where someone present made a statement of fact or expressed an opinion is considered good practice, as it provides a formal record that can help resolve misunderstandings or disputes.

Tip: Sharing meeting minutes is good practice to improve audit results, but it is not the same as sharing initial conclusions.

Escalation

There are two types of escalation occurrences:

1. The first type of escalation occurs during the audit work when the auditor identifies an issue of material risk that is time-sensitive (e.g., a publicly readable S3 storage bucket containing customer PII). Regardless of whether the issue is resolved at the time or whether it is in scope, the auditor would communicate this issue to the client. If it is about material noncompliance with a legal or regulatory matter, it may need to be escalated to the relevant authority. It is strongly recommended to solicit specialist advice before doing this.
2. The second type of escalation occurs after the report's preparation is finished but before publication. It may be linked to escalations during the audit or could be based on a broader concern after the auditor has fully analyzed the findings.

There is normally a clearly defined process once an issue becomes an audit finding, through to acceptance or dispute by the control owner, and then the tracking of activities that flow from that. An auditee that accepts the audit finding logs a response. The control owner might need to provide a statement to audit. The auditor may accept the audit action and schedule a follow-up check.

Final Report

In the final report, unconfirmed findings are removed, but the auditor may add guidance concerning those findings even if they are not present in the final report.

There are three kinds of reports: internal, limited and public. An internal report will cover the prescriptive remediation recommendations with details and action steps (if there are findings). It will have management comments—not for failures, but for circumstances in which controls can be improved (these might be phrased as calls to action for best practices).

The limited distribution report does not disclose detailed technical information that an external party could exploit to cause harm to the organization. As the description suggests, it is given only to a small group of people outside the organization only if legally or contractually required (such as an acquiring bank or a regulator). Otherwise, it is distributed only to a small group internal to the organization.

A public report similar to an attestation of compliance (in PCI language) may be prepared and distributed more widely. It would be classified as a public document with theoretically unlimited distribution. It would not contain specific details of controls or any information that might be sensitive to the organization.

Most often, an auditor will produce an internal report and a public report, or a limited and a public report. It is rare that an auditor would need to produce all three types.

5.10 Chapter 5 Knowledge Check

REVIEW QUESTIONS

1. Identify **ALL** the types of assurance.

A. Audit
B. Attestation
C. Assessment
D. Assertion

2. Mark all of the items that are **UNIQUE** to auditing cloud services and cloud service providers:

A. Knowledge and skills related to cloud technologies and platforms
B. Conducting audits of technology services without access to the assets, controls and facilities that house the operations, infrastructure and platforms of the service provider
C. Defining the audit scope and approach within the context of relevant ISO standards
D. Developing third-party service provider reports in accordance with ISAE 3402 standards

3. Fill in the blank: An integral component of audit execution is reporting and monitoring; hence, audits of public cloud providers require ______________ to ensure audit report responses and corrective action plans are obtained from the auditee.

A. Verbal agreements
B. Description of expectations in audit approach and planning documentation
C. Discussion with auditee point of contact
D. Formal description of expectations in contract terms and conditions

4. One technical difference between the auditing of a cloud computing environment and an on-premises environment is:

A. The audit on the part of the cloud service customer (CSC) focuses on the infrastructure aspect including computing, storage and physical security of data centers.
B. The audit on the part of the cloud service customer (CSC) primarily focuses on platform and workload created on its platform, but not the infrastructure of the cloud service provider (CSP).
C. The audit on the part of the cloud service provider (CSP) not only focuses on its infrastructure and services but also on the platform and workloads of the cloud service customer (CSC).
D. The audit on the part of the cloud service provider (CSP) only focuses on the platform and workloads of the cloud service customer (CSC).

5. How does the auditing plan for cloud governance fit into the existing company plan?

A. Cloud service customer (CSC) policies, process and internal controls that comprise how the organization addresses the handling of cloud arrangements according to used deployment models
B. Cloud services provider (CSP) ability to adjust its service development to the newest technological features is independent of the overall governance chain
C. Contract between a cloud services provider (CSP) and a cloud service customer (CSC), including cloud services provider (CSP) assessments, documentation on a cloud services provider (CSP), internal (i.e., self) and external compliance assessments are not relevant.
D. Contract between a cloud services provider (CSP) and its supplier is excluded from the governance of cloud governance

6. The regulatory environment in which an organization operates affects the scope and extent of the audit. Choose the correct statement(s) below:

   A. FedRAMP has more stringent requirements than CSA Star.
   B. CSA Star has more stringent requirements than FedRAMP.
   C. The more stringent the regulation, the more rigor required of auditors.
   D. The more stringent the regulation, the less rigor required of auditors.

7. Which of the following areas are difficult to evaluate through automated continuous auditing?

   A. Authentication
   B. Documentation quality
   C. Backup and recovery
   D. Vulnerability assessments

Answers on page 292

Chapter 5 ANSWER KEY

Correct answers appear in **bold** font.

1. **A. Auditing is a type of assurance service that includes a systematic process of objectively obtaining and evaluating evidence regarding assertions about certain subject matters. Auditing ascertains the degree of correspondence between the assertions and established criteria, and involves communicating the results to interested users (as stated in 6.2.4 Table 1).**

   B. Attestation engagements are another category of assurance service in which an external auditor is engaged to issue a report on subject matter that is the responsibility of another party, e.g., agreed upon procedures, reporting on controls at a service organization (as stated in 6.2.4 Table 1).

   C. Assessment goes further than an audit, as it involves the determination of actions necessary to make the assessed entity compliant (as stated in 6.2.4 Table 1).

   D. Auditing management assertions are a specific type of assurance service (6.2.4).

2. **A. Auditing the security and privacy of cloud computing requires a combination of a variety of skill sets and expertise including cloud technology skills and expertise.(6.3,2).**

   B. In a public cloud the organization loses control over infrastructure, oversight, audit and enforcement (Unit 7.2.2.2).

   C. Incorrect (6.3.1.1) These standards were developed prior to the evolution of cloud.

   D. Incorrect (6.3.1.3) These standards predate cloud computing).

3. A. Incorrect, verbal agreements are not enforceable.

   B. Incorrect, audit documentation is for cloud customers, not enforceable with cloud providers.

   C. Incorrect, needs to be documented and approved by contract

   D. Correct, A recommended best practice related to reporting and monitoring is to ensure the service agreement with the CSP includes the right to conduct audits, as well as the CSP's responsibility to respond to audit results and agree to corrective action plans. (section 6.4.3.2)

4. A. In cloud computing the cloud service provider (CSP) is responsible for the infrastructure including availability of computing, storage and physical security of data centers.

   B. According to the shared responsibility model cloud service customer (CSC) has to ensure that compliance and security aspects of the platform workload are developed on it.

   C. In cloud computing the cloud service provider (CSP) is responsible for the infrastructure including availability of computing, storage and physical security of data centers therefore the audit on cloud service provider (CSP) focuses only these aspects but not the platform and workload made by the cloud service customer (CSC).

   D. In cloud computing the cloud service provider (CSP) is responsible for the infrastructure including availability of computing, storage and physical security of data centers therefore the audit on cloud service provider (CSP) focuses only these aspects but not the platform and workload made by the cloud service customer (CSC).

5. **A. Cloud service customer (CSC) policies, process and internal controls that comprise how the organization addresses the handling of cloud arrangements according to used deployment models that are part of general best practices in almost every established and reputed organization.**

   B. All efforts within the organization whether it belongs to performance and conformance shall be incorporated into the existing processes.

C. Contract between a cloud services provider (CSP) and a cloud service customer (CSC) including as well as cloud services provider (CSP) assessments, documentation on a cloud services provider (CSP) internal (i.e. self) and external compliance assessments provide clear understanding of handling of certain controls and hence it has to be commenced and evaluated accordingly.

D. Contract between a cloud services provider (CSP) and its supplier must include all the necessary aspects and means to the organization to adopt the new technologies in a way that is digestible and compliant and vaulted in the existing or adopted governance spectrum.

6. **A. FedRAMP, which is a government regulation, is more stringent than CSA Star.**

B. FedRAMP, which is a government regulation, is more stringent than CSA Star.

C. There is a greater expectation on auditors regarding the scope and extent/depth of testing when the regulation is more stringent.

D. There is a greater expectation on auditors regarding the scope and extent/depth of testing when the regulation is more stringent.

7. A. Topics such as backup and recovery, vulnerability assessments, static and dynamic code analysis, logging and filtering, cryptography and authentication are well suited for this approach(6.8.6 Continuous auditing).

B. Procedures and documentation quality are much harder to evaluate through automated continuous auditing (6.8.6 Continuous auditing).

C. Topics such as backup and recovery, vulnerability assessments, static and dynamic code analysis, logging and filtering, cryptography and authentication are well suited for this approach (6.8.6 Continuous auditing).

D. Topics such as backup and recovery, vulnerability assessments, static and dynamic code analysis, logging and filtering, cryptography and authentication are well suited for this approach (6.8.6 Continuous auditing).

Page intentionally left blank

Chapter 6:

Evaluating a Cloud Compliance Program

Evaluating a Cloud Compliance Program

Evaluating a Cloud Compliance Program

6.1 Learning Objectives

After completing this chapter, learners will be able to:

1. Outline audit characteristics, criteria and principles.
2. Describe auditing standards for cloud computing.
3. Define auditing an on-premises environment vs. cloud.
4. Recall differences in cloud services and cloud delivery models.
5. Explain audit building/planning and execution.

6.2 Overview

Chapter 2 focuses on designing and building a cloud compliance program for an organization, mainly from the perspective of a cloud customer, looking at the requirements, variables and perspectives to consider for a holistic view of the organization's compliance exposure.

Chapter 6 enables a learner to evaluate an organization compliance posture and assess if its compliance program is effective.

There are three main components in the cloud compliance posture evaluation, and they relate to the idea of the cloud shared responsibility model and the concepts of responsibility and accountability:

1. How does the organization evaluate if a certain cloud service is fit for purpose? How does it evaluate the effectiveness of the controls implemented by the CSP? What are the contractual guarantees? Ultimately, how does the CSP enable the enterprise's compliance objectives?
2. How does the organization evaluate the effectiveness of the controls it is directly responsible for?
3. How does an organization facilitate continuous compliance?

This chapter focuses primarily on understanding what an organization should do in order to enforce the necessary control and oversight over the cloud service portfolio and to practice due diligence over the CSPs during the business relationship life cycle, from service selection until its potential sunsetting.

The main reference for the evaluation will be the requirements, criteria, objectives and standards addressed in chapter 2.

Note: Evaluating a cloud compliance program is essential to determine the reliability and accuracy of the measures in place to govern the requirements for which the organization is accountable.

6.3 Compliance Program Evaluation Approach

The oversight and monitoring of the CSPs and the cloud service portfolio are done by using the governance tools introduced in chapter 1 (e.g., policies, audits, vulnerability assessments, cloud platform assessment, penetration tests, contracts and third-party certification and attestations), and evaluating mitigating actions to ensure that the risk levels are maintained within the limits the organization is willing to accept. The audit function will play a key role in this evaluation process.

The internal compliance program team that originally worked on building the organization's compliance program will certainly take steps toward measuring its effectiveness. Some actions may include collecting internal policy attestations, control testing, and real-time reporting through security tooling. There may even be cases in large organizations of an entire business function (CISO, Business Controls, business/process owner, etc.) monitoring for compliance by conducting internal audits from a different perspective than the team that built the compliance program.

Note: The terms auditor and assessor can differ in meaning and nuance. In the context of this chapter and throughout this guidance; the words are used interchangeably.

External auditors will be purposely looking for control gaps independently of the company. The external view is to get a snapshot of the current environment and how the compliance program measures at a given point in time. When the audit is complete, the external auditor will provide, if applicable, a list of observations the internal team can use to remediate the issues noted. It is extremely beneficial to perform the remedial action and audit when preparing for a third-party evaluation (e.g., for certification or attestation purposes).

The approaches to an evaluation within a large or a small organization can vary. Within a large organization, there are generally dedicated business units responsible for internal audits. They coordinate with the governance, risk and compliance (GRC) and security operations center (SOC) teams and infrastructure owners to perform routine assessments and evaluations at different levels of the business. Some of these teams include pen testers and SOC analysts, for example. These internal audit teams will work with external auditors that may be invited to do an independent third-party assessment when the company is being held accountable to a certain requirement (e.g., PCI-DSS or SOX).

The organization may work regularly with an auditor to conduct an evaluation when needed. After establishing the relationship with the business, the auditor can understand and channel it into a compliant state, and a better overall security posture, by putting lessons learned and feedback and to work throughout the business. It is much better to have a regular auditor who is already familiar with the organization conduct an evaluation when needed to ensure a smoother and less arduous but auditable process. A new auditor will require complete onboarding to gain a good grasp of the organization from the ground up.

The same concept applies to small companies, except that a smaller company will probably not have a dedicated business function for internal audit. The employees who perform the GRC work will more than likely be the ones working with the auditors performing the evaluations. Those employees might need to take on more responsibilities, such as iterating on the GRC program, testing and remediating controls, monitoring and measuring controls' effectiveness, auditing and evaluating against GRC requirements, and reporting status. By contrast, these tasks are assigned to different parts of the business.

Tip: Depending on the size of the organization, the auditor may need to work with many people rather than just a handful.

6.4 The Governance Perspective

To evaluate their effectiveness, it is important to understand the problems the compliance and governance programs need to solve, and the purpose each one fulfills. The two programs work in tandem. The cloud governance program identifies what external and internal requirements the organization should adhere to, and how to comply with them. The cloud compliance program measures and validates the effectiveness of the controls in place to mitigate the risk of exposure.

6.4.1 Evaluate the Consistency Between Cloud Governance and Cloud Compliance Programs

An organization needs the cloud governance and compliance programs to work in sync. The bidirectional communication channels between the organization and governing bodies, auditors (internal and external), and legal departments, etc., need to align so the organization can stay ahead of new requirements or changes that need to be met (**figure 6.1**).

Figure 6.1—Maintaining Compliance

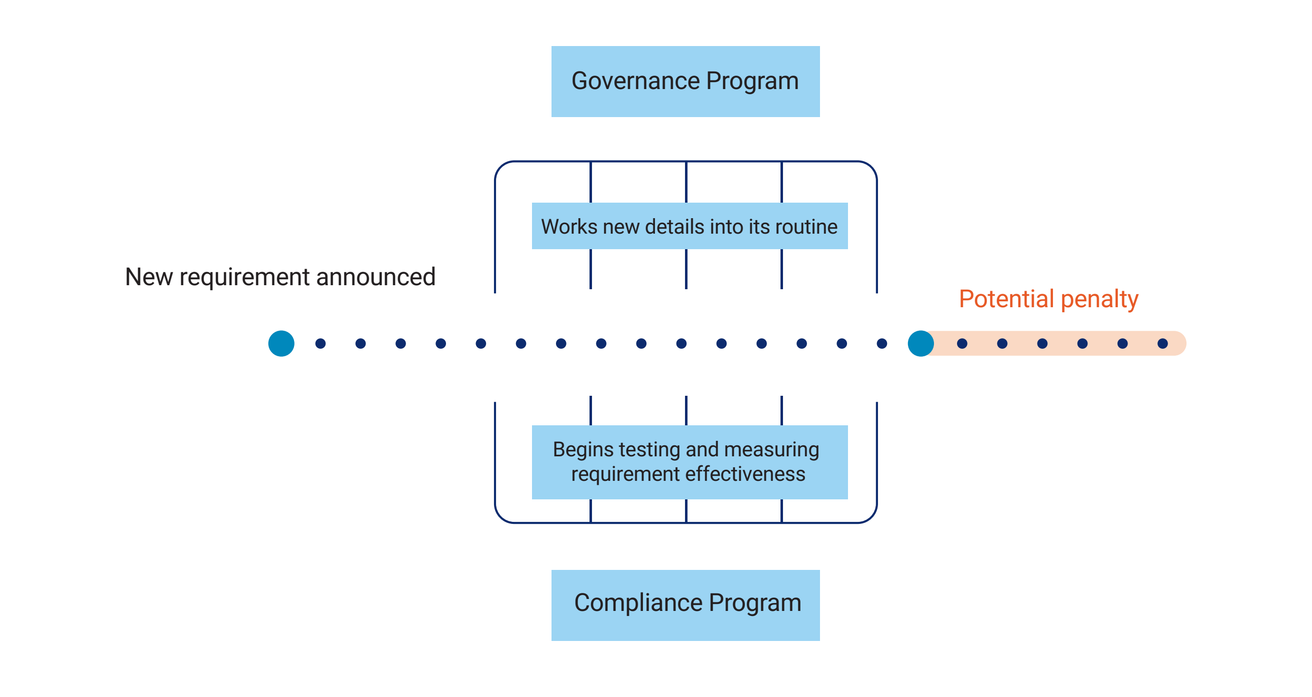
Governance Program
Works new details into its routine
New requirement announced
Potential penalty
Begins testing and measuring requirement effectiveness
Compliance Program

Once a new requirement is announced, there is an adoption period for the organization to become compliant. During this time span, the governance program needs to work the new details into its routine, and the compliance program needs to begin testing and measuring their effectiveness. If the organization does not meet the requirements in time, there may be penalties for not complying. Having a strong communication channel in place to identify these changes before it is too late reflects the organization's maturity.

While working to connect the two programs, the organization should identify any gaps between them and determine what steps are required to ensure the gaps do not cause additional risk. There should be measures in place that constantly check for issues. Many organizations choose to leverage solutions to help find discrepancies and to provide alerts for new requirements or changes. Changes will be needed because new requirements are always emerging as technology evolves. Being able to adapt and drive updates into the organization is the key to success.

6.4.2 Policies Evaluation

An assessor who is tasked with verifying if the program is effectively implemented should be able to address questions related to policies, such as the following:

1. Is the policy correct? Does it reflect the requirement?
2. Is it communicated, implemented and enforced effectively?
3. Does it achieve its objectives?

Some key actions to perform:

1. Perform pre-assessment[248] activities to review the policies and procedures and ensure the objectives of the assessment can be met.
2. Pre-assessment activities may entail performing interviews to understand the real and perceived meaning of the policies/procedures (don't take evidence at face value). Depending on the size of the business, the assessor would need to have open communication channels with different units and stakeholders (to ask questions to see if the responses reflect what the policies/procedures state). For instance, once the human resources (HR) department clarifies how security awareness is enforced and conducts a spot check if the policies/procedures are being carried out, is there an awareness training program in place? Is there an allocated budget? How much time can employees dedicate to security awareness? Is it mandatory? Are the employees' awareness levels verified?
3. Depending on requirements, get preliminary evidence of the policies/procedures and ask questions about who makes the decisions, and if and how policies/procedures are communicated. Request examples of enforcement techniques, etc. This preliminary check will help the assessor to form an initial opinion on how things are run in the enterprise, including strengths and weaknesses. Is the compliance program likely to be just a theoretical exercise, or is it effectively implemented in practice?
4. Start with the formal, planned assessment[249] of the policies/procedures and associated processes (after the previously described actions have been performed).

Define Responsible Parties for Policy Enforcement

An auditor who is evaluating the target organization use of the cloud should become familiar with how it formulated its cloud policy, how decisions were made, and who made them. What are their backgrounds? Do they have cloud skills and experience, and do their decisions reflect the risk appetite of the company?

It is valuable for the auditor to interview the people who formulated the policies and who worked on the organization's earliest cloud adoption. This is useful to understand the pace of policy development set against lessons learned from practical cloud IT projects, e.g., a proof of concept that led to the development of a policy.

Evaluating Cloud Usage Impact on Corporate Policy

During the policy development stage, large organizations with multiple business units may have competing and conflicting cloud objectives or usage models. The auditor needs to determine how the policy owner has reconciled potential differences between different cloud objectives to arrive at a group policy that enables it to operate within its risk appetite, while fully leveraging various cloud services in the way that is aligned with the organization's strategic plan.

The assessor is advised to investigate the organization decision-making process. If there is an enterprise-wide architecture board or committee, cloud policy decisions will be influenced by its decisions. The board or committee will have access to meeting minutes reflecting decisions taken. Good policy design in part involves the stakeholders presenting different views that result in trade-offs and acceptances. The auditor or assessor checks whether major stakeholders agreed with the policy or not. If not, there may be a risk they will act in bad faith or choose not to follow it.

[248] Typical preassessment activities that measure readiness include: a) reviewing documented information, b) evaluating site-specific conditions and undertaking discussions with personnel to determine preparedness for a formal planned assessment, c) reviewing status and understanding regarding applicable requirements.

[249] The purpose of the formal audit is to evaluate the implementation, including effectiveness, of the policies and procedures.

Policy Violation

The organization must define what constitutes a policy violation (see chapter 2) and be able to identify when one has occurred. The question for an assessor to determine: Has the organization defined in very clear language the conditions for policy violations? Does everyone connected with the organization know what they are? Do they know what they are allowed to do with the cloud service, or do they need to check the policy? Is there a clear escalation path in case of policy violation?

By definition, insiders are the ones who typically breach policies since external attackers aren't subject to them. There are three groups of insiders: employees, contractors and external consultants. The level of access they have, or the visibility of the policy violation, determines its materiality.

A policy violation may be trivial. It could be accidental, which is the likeliest scenario, or deliberate. The assessor needs to ask what fail-safes the organization has in place against any violations, e.g., how easy would it be for someone to accidentally delete a cloud account or mistakenly leak data?

Tip: From the practitioner's point of view, the auditor considers the account design as much more critical when setting up a cloud environment. In a large organization, a master account is often given to divisional technology leaders, while suborganizations create cloud user accounts. The account design model is driven by billing and security. Auditors usually encounter approved cloud accounts, but employees also may use cloud accounts that the organization has not formally sanctioned. These kinds of cloud workloads are known as shadow IT—paid for by employees themselves (or a consultant/vendor) and not falling under normal governance structures. Because these types of accounts can represent a security risk, it is vital for an auditor to ascertain that no account falls outside the scope of the audit.

Cloud-Centric Policies and Automation

One of the most radical changes the cloud introduces is codification into software of all the components of a service architecture (virtualization of servers, networks, etc.). This is well summarized by the idea of Infrastructure as Code (IaC).[250] The idea of everything as a code stems from the need to keep pace with the changes and scale of cloud services. Clearly, security needs to adapt and move toward automation, and therefore needs to change policies and the evaluation process (i.e., compliance automation).

In a mature cloud environment, policies are often translated into code as opposed to just being text documents. Likewise, the controls need to verify that the policy is effectively implemented. So, policies and controls are directly implemented in the system in a (semi) automated fashion.

From the cloud compliance program perspective, this means that most of the evaluation on the suitability of policy to effectively manage the risk and align with the governance goal is done via the analysis of audit trails (e.g., logs). Logging of activities takes place alongside the development, delivery and integration pipelines to determine who has done what, who has approved what, when something has been tested, when a certain change has occurred, etc. The logs are not the only evidence the assessor will rely upon, but their role is much more central than before.

See chapter 8 for more details about DevOps and continuous compliance.

[250] Schults, C.; "What Is Infrastructure as Code? How It Works, Best Practices, Tutorials," Stackify, 5 September 2019, https://stackify.com/what-is-infrastructure-as-code-how-it-works-best-practices-tutorials/

6.4.3 Evaluate How the Program Tracks Controls Owned by the Internal Organization

An important concept for any cloud compliance program evaluation—and for cloud auditing more generally—is that the audit of an organization should not be qualified as failed because the organization does not have a cloud policy. Instead, an assessor needs to determine the following:

- The organization has identified the security risk associated with cloud usage.
- It has evaluated how to manage the risk.
- It has at least a work-in-progress set of control objectives that respond to the risk.

A helpful way to ensure that internal controls are aligned with the organizational requirements and policies is to identify mechanisms within and outside the organization that determine the following:

- The organization risk
- How to prioritize the risk
- The controls necessary to mitigate the risk

See chapters 1 and 2 for more information on the process of identifying and assessing risk and mitigating it through controls. From the assessor's perspective, it is important to answer these questions:

- Who designs the controls?
- Who implements the controls?
- Who operates or executes them?
- Who owns them? Who is the control owner, i.e., the party responsible for making sure the control operates as defined and accountable for making sure the control is enforced? It may be helpful to think in terms of the RACI model: Who is informed about the health of the control (see section 1.2.4)?

With the cloud's shared responsibility model, the assessor or auditor must establish how the organization has approached the process of designing controls against its own risk management framework and policies. It is the auditor's responsibility to call out where those policies fall short of industry best practices. The auditor or assessor's role is to review not only the design process but also the governance of decision-making, the management support, the risk identification and mitigation (and risk acceptance, if appropriate). The auditor or assessor should consider mitigating or compensating controls and recommend potential improvements the organization must consider if a control fails. There must be documentation reflecting decisions taken, particularly the decision on whether to proceed with adopting and implementing a cloud service on or off premises. In addition to the use of cloud services, explicit documentation must include where data are stored and processed.

Tip: The risk owner is not the technology department but is the one identified through governance to be the data owner. This is usually a business function.

6.4.4 Third-Party Risk

As organizations increase their dependency on third-party services to underpin and deliver key business processes, greater competency is needed to understand how the use of outsourced distributed technologies shapes its risk profile. To gain the benefit of using these services, the organization technologists and its legal, privacy, compliance and risk practitioners need to be able to identify, manage and respond to cloud risk so that back-office risk management and control costs do not eclipse the benefits of using them.

The assessor needs to evaluate the organization decision-making process for its selection of cloud services based on three levels of analysis:

1. **Sourcing**—How does the organization source its business partners? How does it onboard them in a way that is commensurate to the risk? What is the due diligence process? What service level agreements are in place? Are there assurance documents to check?
2. **Risk evaluation**—Consider the individual services that an organization consumes from the CSP—e.g., AWS and Azure provide a broad catalog comprising hundreds of services. What are the risks and controls of those services (which will vary widely in the maturity of what they expose to their customers)? The CSP may have two levels of assurance documents for its services—i.e., one set that it makes available to everyone, and another enhanced set of assurance documents available to organizations that subscribe to its compliance program.
3. **Risk mitigation**—How does the organization use the services it has bought from the CSP? What controls are part of the service, and what controls does (or must) the organization (need to) place on the account? The auditor must assess the following:
   - What are the skill sets of the people the organization has given permission to (paying particular attention to the administrators)?
   - Are the security controls equivalent to what it has on premises?
   - Do the organization security and administration people have the skills to maintain those controls and understand when they are not working?

There are three types of controls: preventive, detective and corrective. The assessor considers not only whether they are the right controls, designed in the right way and applied in the right places, but also how the organization analyzes the output from those controls to identify threats.

Tip: The controls of an organization must adapt and evolve depending on changes in the threat landscape. A question for the assessor to ask is this: Over a given time period, how frequently is the organization reflecting the latest threats its organization or its clients are facing? Has that led to any changes in the choice and deployment of controls?

CSP and Cloud Service Inventory

A typical cloud implementation scenario is a hybrid cloud model that allows the cloud customer to integrate several public cloud services, ranging from IaaS to SaaS, into its existing on-prem infrastructure or private cloud. The number of vendors involved in such a hybrid IT architecture might vary from tens to thousands (depending on the size and needs of the organization). As the asset inventory is a fundamental tool of any security management program, so too is the inventory of cloud services and providers. This is a much-needed tool to maintain a holistic view and keep control over the cloud service architecture.

A cloud customer faces risks related to hidden interdependencies between cloud services, and the root cause for them is lack of visibility over the cloud service supply chain and architecture, either because of lack of transparency from some CSPs about their chain of subcontractors or because of the well-known phenomenon of shadow IT.

The shadow cloud service typically avoids or evades appropriate risk management and due diligence assessments. This creates weakness in the organization’s compliance posture, since by definition the shadow cloud service is not identified or evaluated through the normal governance, policies and standards, and procedures and protocols established by the organization.

From the assessor’s perspective, it is important to verify if the organization has policies, processes and tools in place to identify cloud services that operate outside the boundaries of the company governance. Besides introducing a policy that explicitly forbids unapproved services, cloud customers might need to implement service discovery tools (e.g., cloud access and security brokers, also known as cloud security gateways).

Key Factors to Consider

The key factors for assessing whether the cloud service meets the customer's expectations include the following:

- **Regulations and standards specific to a region**—If the customer needs to comply with a certain privacy and data protection regulation (e.g., the EU GDPR, Japan APPI or California CCPA), does the cloud provider meet the necessary conditions for international data transfer?
- **Industry experience and reputation**—Does the cloud provider have experience in dealing with customers operating in the same business sector or geographies the enterprise operates in? Does it have a dedicated offering for a certain business sector?
- **Security features**—Does the CSP allow for customer-managed encryption keys (bring your own key, or BYOK)? Or more generally does it implement the CSA CCM controls?
- **Proof of compliance**—What does the vendor have in terms of attestations/certifications?
- **Audit capabilities**—How does the CSP enable the organization's auditing requirements? Does the contract include the right to audit? What are the terms and limitations? What information does the CSP share with the customer?
- **Financial stability and outlook**—Is the CSP reliable from the financial perspective? Does it provide the necessary level of assurance that it will be able to stay in business in the long run?
- **CSP's supply chain management and underpinning contracts**—While the organization may be negotiating with a sales office in Europe, the CSP may be legally localized in the US or in China where the bulk of its operations are based, and it may also have storage systems outsourced in China. This may create issues, depending on the organization's regulatory requirements.

Region-Specific Privacy and Security Regulations

When assessing a cloud service provider compatibility, the customer needs to identify its responsibility for compliance with specific standards and regulations (e.g., PCI/GDPR/HIPAA). Privacy regulations often drive compatibility from a governmental standpoint, especially when it comes to complying with local, state and regional laws (CCPA, GDPR, etc.). Given those requirements, do the CSP or the cloud provider headquarters meet the necessary geographical location and data center location requirements for each of those laws/standards for consideration of data residency?

Industry Experience and Reputation

When assessing the compatibility of a service with requirements, it is important to consider whether the cloud provider has industry-specific experience. This is important both from the perspective of the overall suitability of the CSP for inclusion in the list of approved providers, but also from the specific audit standpoint. Dealing with a CSP that has experience in the particular business sector will provide additional assurance that it can satisfy the industry requirements and provide the much-needed supporting evidence during an audit.

Even if the CSP has a better product than competitors, it may not have the necessary levels of customer service. Consider whether it is able and prepared to provide the necessary support to the customer during scheduled or unscheduled audits. During an audit, will it be able to support the customer and provide visibility into the provider's environment?

> **Tip:** In the PCI world, it is typical to get a responsibility matrix (RACI chart) for the cloud provider. This examines the whole PCI standard and all controls, and maps whether each is a provider control, merchant control or shared control. This exercise allows the customer to see what the merchant (customer) is responsible for. It also helps to manage risk, because it allows a customer to see where it has transferred risk to the provider.

Security Features

Security features are among the most important variables to consider in evaluating a cloud service—both the security features the CSP is responsible for and the capabilities (typically APIs) offered to the customer so that it can plug in additional services (either custom-built or bought from a third-party security-as-a-service [SecaaS] provider).

Assessing a CSP or a cloud service against its security capabilities and features typically follows best practices and standards. CSA recommends the use of its open-source tools CCM and CAIQ.

Proof of Compliance and Right to Audit

The cloud places limits on the kinds of assessments an auditor can carry out. Many CSPs do not permit penetration tests or vulnerability scans, but large providers publish a policy around their own security controls and checks. A cloud contract typically does not include the right to audit. It is worth noting that policies usually evolve over time and allow more testing as they mature.

The auditor will rely primarily on the CSP statements. If a provider offers a summary of the results of the assessments it had carried out by a third party, the auditor will want to know if the scope of those assessments is material and sufficient to the services the auditee is consuming. This scope is critical, as it refers to attestations of security. Effectively, the auditor wants to establish the relevance, quality and depth of any third-party security assessments: Who performed the assessment (what are the skills, expertise and credentials of the testers)? How much time did they apply to the test? Were they operating under any constraints for the test?

The auditor may engage with the cloud provider's security team. The nature and speed of its response to security questions could contribute to the auditor's qualitative judgment. The auditor also aims to establish whether the auditee's person responsible for security will simply accept evidence of security testing from the CSP, or whether the auditee will follow up with questions to establish more clarity regarding security controls.

Depending on the level of assurance the cloud customer requires, it is important to understand if the CSP has a proof of compliance that is validated by a third party. Several cloud-relevant security certifications and attestations exist, e.g., CSA STAR Certification and STAR Attestation, ISO/IEC 27001 (when supplemented by the ISO/IEC 27017 and 27018 controls), SOC 2.

> **Note:** Although certification and attestation are key pieces of information (sometimes even a minimum condition) for understanding a CSP security and compliance posture, they are not necessarily sufficient to satisfy the level of assurance the customer requires. For example, in the financial sector, the regulators require the financial services provider to evaluate the risk associated with their use of cloud, not just based on the certification and attestation held by the CSPs. In those cases, the auditee (and the auditor) should require additional information to verify the provider security and compliance standing.

Vendor Lock-In

In economics, vendor lock-in, also known as proprietary lock-in or customer lock-in, makes a customer dependent on a vendor for products and services, unable to use another vendor without substantial switching costs. With regard to cloud services, vendor lock-in may be a consequence of proprietary technology, architecture or data formats specific to a single vendor, prohibiting or significantly delaying a possible data migration to a new CSP.

Financial Stability and Outlook

Besides examining technical capabilities and contractual guarantees, it is important when evaluating a third-party service to understand if the provider will be able to offer the service in the long term, to avoid so-called vendor lock-out.[251] It is important to consider its financial stability and outlook. Financial information is typically available to everyone for consultation, and it is an important job for the assessor to consider it when evaluating the CSP's compliance outlook.

Corrective Action Requirements

The assessor should define at a high level what responses the cloud customer wants from its provider in the event of a security incident. In its contract, the customer may require a CSP to give notice of a security breach or provide information on a foreseeable event that is likely to impact its ability to service the customer's workloads.

The assessor needs to ensure that the shared responsibility model is reflected in the SLAs, contracts or agreements, so each party understands what actions they're responsible for when it comes to identifying, correcting and reporting deficiencies or risk events.

6.5 Legal, Regulatory and Standards Perspectives

When designing and building a compliance program, it is important to consider what national, regional, or international laws and regulations apply, and whether there are applicable statutory or industry requirements, to gain an understanding of the contractual conditions and obligations arising from cloud agreements.

6.5.1 Evaluate Alignment With Legal, Regulatory and Contractual Requirements

Once the legal and regulatory requirements are defined, the organization must evaluate the following:

- Whether the measures implemented are suitable to satisfy the requirements imposed by the applicable laws and regulations
- How to monitor the legal and regulatory landscape to ensure that changes in the applicable legal framework are promptly reflected in the compliance program
- How the changes in the service architecture could create a misalignment within the compliance program
- How the contractual terms impact the compliance program

None of the actions mentioned previously are different from those required in noncloud environments. The main difference when dealing with a cloud compliance program is the pace of change of the service architecture and service portfolio.

The fact that cloud computing services are based on complex supply chains and that there is a high level of dependency on third parties places additional emphasis on the effectiveness of the company vendor management program. When evaluating the compliance program, it is crucial to determine how any changes in cloud contract terms might impact the enterprise's compliance posture.

Some issues to be evaluated and key questions to be asked follow:

- What are the contractual terms to be monitored?
- What are the technical terms to be monitored?

[251] This occurs when the customer is unable to recover or access its own data because the cloud provider has discontinued the service, either by entering bankruptcy or otherwise leaving the market.

- How does the compliance program monitor those changes?
- What are the procedures to react to changes?
- If the risk generated by the changes cannot be mitigated, what is the exit strategy?
- Is a temporary loss of legal compliance affordable for the organization?

In terms of contractual and technical changes, it is important to review the following:

- Terms and conditions or master services agreement
- Data processing or privacy addendum (e.g., the GDPR Addendum
- Service level agreement
- Acceptable use policy
- License agreement
- Supply chain structure
- Escrow agreements

6.5.2 Contract Analysis

An assessor who reviews contracts should attempt to find out how the customer has addressed its then-current needs and any possible gap between its previous terms (if the service was previously delivered in-house) and the ones available via the CSP.

Relationships with suppliers are more than the sum of a contract. The contractual agreements and supporting documents help CSPs mitigate or transfer some risk associated with offering services to unscreened third parties. They provide CSCs with legally binding assurances of how the CSP will handle their accounts and data.

Beyond the contract terms, the assessor should evaluate the practical effect of the contract, such as the CSP's willingness and ability to provide information about security or privacy issues. An assessor should flag behaviors that have led to, or could reasonably lead to, poor security outcomes, e.g., consistent delays between a customer's request for information related to operational security and a suitable response from a CSP.

The assessor should also seek to establish whether the organization is aware when the contract terms change, such as when SLA terms change. This monitoring could be performed either through a procurement manager or third-party services that monitor changes in the terms of cloud contracts.

Subcontracting and SLAs

The most effective way to ensure the CSP complies with the agreed-upon service is to document the requirements and responsibilities specified in the customer contract and SLA agreement.

When evaluating cloud SLAs and related third-party reports, the auditor should compare them to the organization's internal control framework, validate whether it meets the requirements, and whether it meets industry best practices (CCM, for example).

Normally, the evaluation involves mapping controls from the organization's internal policy and standards to the cloud provider's SLA. For example, the NIST Cybersecurity Framework[252] provides a common language for intermediate mapping. The auditor also evaluates the cloud SLA through a BCP/DR lens to identify the recovery timeline for the service and the potential impact, or whether the organization accepts the risk. This may result in

[252] *Op cit* NIST, "Cybersecurity Framework"

modifying internal SLAs so the organization develops a resilient architecture that considers availability zones or multiple cloud providers.

A CSP is responsible for delivering the cloud service, but the cloud customer is ultimately accountable for the use of that service and the resulting implications—for example, if a CSP doesn't allow a customer to comply with a certain regulation, such as GDPR, it is the customer that is ultimately accountable to the regulator (**figure 6.2**). While customers will want to reap specific business rewards (e.g., faster time-to-market for new services) from using the cloud service, all customers should first seek to protect their enterprises, and in some sectors their executives, from legal, compliance and operational risk. Since any customers can be materially impacted by poor subcontracting choices made by their CSPs, legally binding disclosure and notification rules covering subcontracting arrangements may be required.

CSPs offer convenient APIs to encourage customers and other market participants to build upon and extend their platforms. Viewed through a productivity lens, this empowers developers to rapidly develop and deliver scalable business services by weaving together services from multiple CSPs. Looked at through a governance lens, a customer that consumes this service expands the cloud supply chain from one to many overnight. Without the right controls in place to identify and gain assurance over third-party services, the customer will lose line of sight over who is responsible for what and could be in breach if personal or regulated data is involved.

Figure 6.2—Subcontracting and Service Level Agreements (SLAs)

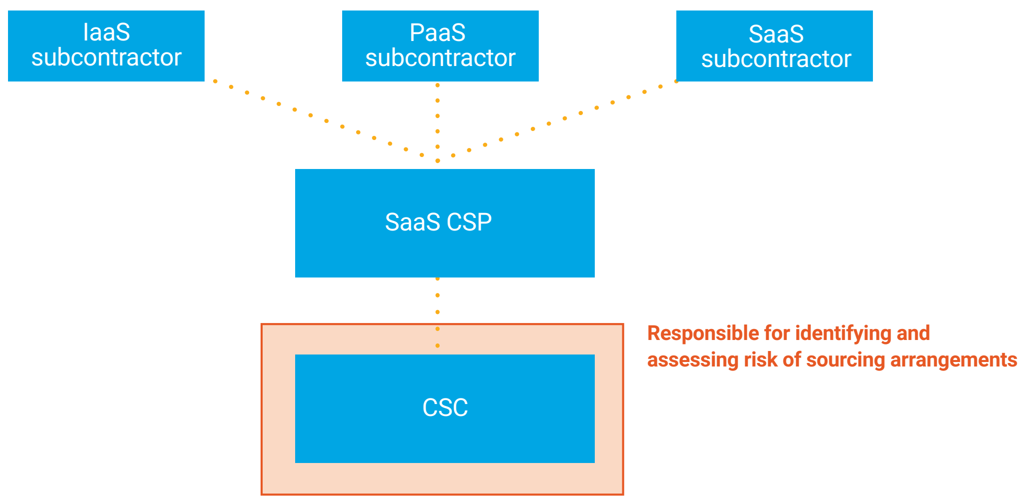
IaaS subcontractor
PaaS subcontractor
SaaS subcontractor
SaaS CSP
CSC
Responsible for identifying and assessing risk of sourcing arrangements

Example: A CSC may enter into a contractual agreement with a SaaS startup that is built on an IaaS infrastructure. In turn, the startup SLA may refer to the IaaS SLA and customer agreement. If an outage or data breach occurs, who is responsible? Clearly, the CSC will look to the SaaS provider, because, in this context, the customer has no direct relationship with the IaaS provider. What happens when the startup decides to switch overnight indexing of customer data to another provider with a lower compute cost? Assume the startup deems the switch of supplier as administrative and chooses not to notify its customers of the change in subcontracting arrangements, as per its standard terms and conditions.

As it turns out, the new CSP is a compute broker that owns zero computers. Instead, it runs a cloud-hosted matching engine to dispatch clients' compute jobs to tiered CPU pools around the world for processing. While anyone can offer idle CPU pools to the public tier, only registered companies can offer gold tier CPU pools that attract premium pricing. Dark Pool is a CPU pool run by an organized crime gang. It is powered by unsuspecting victims' CPUs, ensnared by its malware of the month. Toward the end of each month, Dark Pool operators adjust their pricing to ensure that they are always the cheapest gold tier pool. This is a data magnet for costly month-end billing runs and thus attracts the largest jobs with the strictest completion deadlines from the compute broker. The personal data received in Dark Pool is quickly identified as such by a predatory algorithm that correlates and classifies identities. These stolen identities are subsequently matched to social media profiles, ranked by wealth, and sold in batches to smaller players on an underground data market. Downstream criminal gangs consume these data via a cloud service and apply their own special tradecraft to commit fraud and stalk their victims.

Regardless of outcome, the CSC is responsible for identifying and assessing the risk of its sourcing arrangements on a continuous basis. (Large financial institutions, for example, have a cloud governance committee that normally needs to review and approve material changes involved in moving the organization's data to the cloud or changing hosting providers.)

In the previously referenced scenario, the SaaS startup breached its own contractual terms because it failed to notify its customers of the new subcontractor. If the notification had occurred, the CSC could have challenged the CSP to provide evidence of its own supply chain due diligence of the subcontractor's controls. If the customer assessed that the new subcontractor failed to meet minimum standards or exceeded the risk appetite in other ways, the customer would have an opportunity to respond—to manage the risk. For example, the customer might seek to opt out of the new arrangement to protect its workloads from exposure to the new subcontractor. The CSP might reciprocate with a price change or seek increased legal liability from the SaaS provider. However, such an accommodation would have to make financial sense (in the event of a CSP failure it may go bankrupt) and be credible—that is, it cannot compensate for serious control failures.

The customer, in good practice or based on applicable legal requirements, needs to inform data regulators, sector regulators, and end users of the data breach. The implications go further:

- If relevant regulators deem the client failed to execute its responsibilities related to data protection and operational risk, exposing its customers to (potential) harm, it could incur regulatory sanctions and suffer reputational damage.
- An impacted CSC that had previously taken out cybersecurity insurance to transfer its risk may find that the terms of coverage may preclude payouts if key sourcing controls were not present or failed.
- CSCs may themselves be in breach of binding legal agreements with their own customers regarding compliance or subcontracting.

Auditors must consider contractual terms and supplier management in light of applicable regulatory requirements and other relevant contracts that bind the customer.

Even if a customer is not in a regulated sector, if it processes personal data it is likely subject to data protection requirements (e.g., GDPR). Consequently, many cloud customers will be obligated to identify and manage supply chain risk. This can be challenging with cloud services. However, some practical techniques that can be employed by cloud customers and reviewed by auditors include the following:

- Interviews with pertinent staff in the organization (What do key decision makers know?)
- Reviewing CSP contracts and other materials for disclosure of subcontracting relationships (What disclosures have CSPs made?)
- Examining invoices, corporate card spend and billing arrangements to reveal spend with CSPs, possibly exposing use of services that augment other CSP services, e.g., a CSC having to pay multiple distinct parties associated with the delivery of one user experience (What services are they paying for?)
- Reviewing cloud tenant configurations for third-party API integrations that provide further clues to the supply chain (What third-party services have they permissioned?)
- Analyzing the organization web proxy logs for evidence of cloud usage, including ancillary services that augment or plug into core cloud services the CSC uses (What services do their people use?)

Supply chain relationships at times may be unintuitive or simply lack transparency. In extreme cases, as with the preceding example, relationships may be obscured, intentionally hidden, or obfuscated by financial or technical means.

Auditors should verify that their clients have clear contractual line of sight over their cloud supply chains, that their CSP contracts address ongoing subcontractor disclosures, and that there are no conflicts or discrepancies related to subcontracting across agreements with key stakeholders, such as regulators and their own customers.

Remedies

Holding each party accountable for its end of the deal is generally done through adding consequences tied to noncompliant actions. If contractually, the provider is to provide 98-percent uptime and the customer experiences frequent outages or downtime, the customer will need to hold the provider responsible for not providing the level of service as agreed. Without having proper documented violations, the penalties are not enforceable. Examples of remedies/penalties:

- **Financial penalties**—A portion of the contracted payment is reimbursed to the customer for damages.
- **Service credits**—Funds are transferred into a form of credit or future work granted to the customer.
- **License extension or support**—The vendor extends the license term or level of additional support (i.e., development or maintenance).

It is important to have SLA noncompliance clauses as well. If the managed security service provider (MSSP) does not notify the customer of a breach, who is going to cover those costs after the fact?

6.5.3 Evaluating How the Cloud Service Portfolio Maintains Adherence With the Legal and Regulatory Requirements Over Time

The customer should have a program and process in place to validate—at least on a yearly basis, and every time there's a change in service terms and conditions—that the solutions an organization uses are aligned with the requirements it needs to meet. The customer will need to have a vendor management program, an approved software list explaining the business need, and proof the vendor is complying with the requirements.

Additionally, the statement of service or SLA agreement needs to be revised and updated to make sure no drift occurs with the level of service provided. If the business model has changed internally, the agreement with the provider may need to be updated.

Example: A provider may prepare a white paper that explains specifics about controls available. The white paper may not go into enough detail and just highlight key focus areas of service as a marketing approach to draw in the customer. The customer will need to contact the provider and get more specific information to see if the CSP is able to meet its specific business needs.

The customer also needs to make sure that the documentation supporting its controls is up to date. Business needs might change each year, so the customer will need to follow up with the provider and ensure it is doing its due diligence and providing accurate materials, especially when operating models and business needs change.

6.5.4 Evaluating the Compliance Program Through Data Breaches

The Top Threat Analysis Methodology is a tool to help auditors and risk managers identify governance and compliance program failures, determine necessary mitigations, and understand how to interpret past incidents and breaches as telltales for possible future noncompliance. This approach is necessary, because, despite good intentions, security incidents unavoidably will occur.

A typical compliance program may have multiple incident categories that fall within the scope of regulatory reporting. However, from the cloud standpoint there are two main incident categories to consider:

- **Data breaches affecting personal data**—Included in privacy laws and regulations such as the EU GDPR or the California Consumer Privacy Act (CCPA)
- **Security breaches having an impact on national critical infrastructure**—Such as the European NIS Directive

If a data breach occurs, it is important to analyze it promptly to understand the nature of the incident, whether it formally constitutes a data breach as defined by applicable regulations, whether it is a reportable breach, and what to do next. In many cases, organizations need to acknowledge and respond to the incident within a prescribed time frame. For example, the EU GDPR includes the obligation for data processors (e.g., the IaaS provider) to notify the data controller (e.g., the customer) without undue delay after becoming aware of a personal data breach. Also, the data controller is obligated to notify to the competent supervisory authority no later than 72 hours after discovering a personal data breach.

In the flurry of activity following discovery of an incident, the data owner or data custodian should collect and process information to help answer the following questions:

- What happened?
- What type of data is affected?
- Can the data leak be stopped?
- How can we mitigate the incident?
- Do we have insurance?
- Should we call law enforcement?
- What do we tell them about this, and how do we do it?
- Which laws regulate this circumstance?
- Even if the breach of security would not be reportable, should the organization notify it anyway?

The terms and conditions of processes to follow should be documented in the contract and SLA. The assessor will need to check that the compliance program has provisions to verify that all required contractual terms are in place, that appropriate incident response controls are correctly designed and implemented, and that there are procedures to leverage lessons learned.

6.5.5 Evaluate Alignment Over Time With Standards

The standards and frameworks of reference are selected in the compliance program design phase, so the assessor's duty is to make sure that requirements of the chosen standards are correctly implemented and maintained over time. Like legal and regulatory requirements, the cloud compliance program should have procedures in place to monitor the evolution of the standards of reference, and procedures to verify that a change in the service architecture, security program or information management system would not impact the company alignment with standards.

If an organization decides to use ISO/IEC 27001 or CSA CCM as a reference for its compliance program, it should monitor any updates to those standards. It should also be aware if adopting a DevSecOps approach, for instance, might create temporary or long-term misalignment with the standard.

Very often third-party certification and attestation confirm to the outside world that the organization is indeed following the specifications and rules of the standards.

Evaluate How the Program Maps to Multiple Control Standards

Frameworks such as CSA, NIST and ISO can offer a harmonized set of controls that overlap between frameworks. However, a company may still need to track controls compliance through multiple frameworks in some instances. With cloud services, most companies may not rely on one framework and standard but rather may use a combination of multiple frameworks, standards and vendor reports to map controls.

Ensuring compliance with multiple frameworks and risk models can take time and consume a lot of resources. Companies need to decide if they actually need multiple frameworks—and if not, determine the best control set and framework for the organization. There is typically not a one-size-fits-all solution, but there may be ways to reduce the number of frameworks and simplify, such as CCM.

Example: Company X uses a harmonized set of controls that overlap different frameworks. In its GRC package, a user can track compliance with multiple frameworks specifically by tracking controls across different frameworks as well.

6.6 Risk Perspective

Risk perspective involves evaluating how the program has a comprehensive perspective on sources of risk and evaluating the alignment between the compliance objectives and risk appetite.

6.6.1 Evaluate How the Program Has a Comprehensive Perspective on Sources of Risk

After establishing the organization risk appetite and risk tolerance (see chapters 1 and 2 for details), the customer needs to document its own responsibilities and the responsibilities that require CSP compliance.

The auditor needs to make sure that the compliance program recognizes all sources of potential risk. That means anticipating things the organization would not be able to account for, such as breaking news or published articles. The organization needs to be aware of changing standards, nondocumented interpretations and unexpected vendor noncompliance. A customer may ask to have an assessment done, and the provider might not be able to do so. The provider's noncompliance would pose a risk to the program.

How Does This Get Done?

In partnership, and with a view of the shared responsibility model, the customer and CSP need to initially understand each other's requirements. During the provider and product assessment phase of the project, the customer should already know its expectations for its service provider. If it has this knowledge ahead of time, the customer can simply ask the correct questions to ensure the service provider will not have compliance issues.

Knowing that the organization must have its data stored in different regions may require narrowing down the choice of providers, for a variety of reasons, including legal restrictions on the transfer of personal information across borders. Currently, more than 120 countries[253] have adopted a privacy or data protection law that prohibits the transfer of personal information out of their territory unless the country of destination offers a substantially similar level of protection for that personal information—a concept often summarized as offering adequate protection. In the context of cloud computing services, the privacy laws of the country where the cloud consumer is located usually define what constitutes adequate protection.

[253] UNCTAD | Prosperity for all, "Data Protection and Privacy Legislation Worldwide," https://unctad.org/en/Pages/DTL/STI_and_ICTs/ICT4D-Legislation/eCom-Data-Protection-Laws.aspx

Most privacy laws in effect worldwide provide that adequate protection may be evidenced in a number of ways, such as: an adequacy decision made by the country's highest authority in the matter of data protection; appropriate safeguards provided to the cloud consumer, such as in the form of contracts with special provisions, or binding corporate rules; approved codes of conduct, approved certification mechanisms; explicit consent of the individual; or derogations as defined in the country's privacy law—for example, the fact the transfer is necessary for the performance of a contract concluded in the interest of the data subject whose personal information is to be transferred. In addition, privacy laws commonly require that the data subject have enforceable rights and legal remedies in the destination country, under these mechanisms. It is up to the data exporter to evaluate whether the data importer can provide evidence of such conditions.

6.6.2 Evaluate the Alignment Between the Compliance Objectives and Risk Appetite

This evaluation takes place at the business level, because it will determine how much risk is acceptable for the organization and what it is willing to lose. Based on that evaluation, the compliance team will align compliance with the controls being implemented to remediate risk. This may cause the control base to scale and change, depending on whether the risk appetite changes.

The alignment between compliance objectives and risk appetite is accomplished by leveraging tools such as dashboards or records of vulnerability information. Those tools are used to make information available to the relevant teams, and to connect with other important tools, such as GRC tools. For example, where vulnerability information shows misalignment between compliance objectives and risk appetite, compliance and risk teams could use the data to assess the potential impact of the vulnerability, while the dev and operational teams would mitigate the vulnerability through patching.

Examples: 1. An organization API security policy mandates that internally developed APIs granting access to sensitive data use OAuth2 for authorization. If there is a need for an exception to this policy, it would have to be approved by the risk committee (or risk owner) and justified and communicated to the compliance team. This would reduce the risk of noncompliance. The process of requesting a risk exception should identify possible compensating controls that could be implemented in this situation.

2. Some companies decide to patch every quarter. They may choose to release all noncritical patches on a quarterly basis and address critical patches when they can get to them. This decision is based on the company risk appetite and is not aligned with its compliance objectives. However, once this decision is made, the organization must be willing to accept this cadence.

Evaluating the alignment between compliance objectives and risk appetite is in large part about tooling. It all comes back to level of oversight and assurance that controls are in place and being reported on. Reporting is feedback; if there is no feedback, then no one is being held accountable for knowing if controls are set up or how they should be set up. Most of this work can be automated with the correct tools in place.

Tooling for Evaluating and Tracking Risk

The company will need to use a risk management tool for its risk database. Depending on the level of complexity and size of the organization, evaluating a cloud compliance program may not be possible entirely through spreadsheets or custom-made web forms; it might require a dedicated application that can scale.

It is advantageous for the cloud customer (and the assessors) to have tooling that allows them to compare the risk appetite and compliance posture, and their alignment over time (e.g., year over year), to understand how and why risk changed, and to understand the company's current risk posture. This might help the organization in improving the compliance program and the reactive actions in case of noncompliance.

Some key points to keep in mind:

- Is the risk register tied to compliance?
- Is the company tying it to the security operations center (SOC)?
- Is the company seeing incidents all the way through to remediation? (For this purpose, the Top Threat Analysis Methodology introduced in chapter 4 could be a useful resource.)

For a company to be successful, the best-case scenario is having real-time compliance and security measures in place. Having a continuous lifecycle that integrates everything from the inventory of assets down to the people working on them is critical to align threat intelligence and incident response. With a single-touch vulnerability management system deployed across the company infrastructure, remediating incidents becomes much easier. The idea is to view the management console that is monitoring for threat intelligence and once an issue arises, take action across all assets within the system to slow the spread or remediate entirely. For example, if there is a known vulnerability that is targeting SSH, the single-touch system can send a script to all the servers to close the SSH ports. However, this single-touch setup may introduce new risks, such as availability risks.

6.6.3 Risk Heat Map

It is importance to use cloud service inventory as a tool to maintain a clear picture of the business partners and a necessary foundation to evaluate the risk connected with each service. If an organization uses more than one cloud service, and if there is a limited budget for security and compliance, it is important to be able to define priorities.

An important intended output of the risk assessment is a heat map. As shown in **figure 6.3**, the heat map can be a good tool to provide a consolidated view of the risk ratings of all cloud providers across the organization.

This input will foster the development of the annual assessment/audit plan based on the following:

- The business unit/function that owns the relationship with the cloud provider
- The type of audit or examination based upon the risk ratings
- Whether the cloud provider warrants continuous monitoring based on risk rating or regulatory requirements
- Whether the cloud provider warrants limited reviews or assessments due to a relatively low level of perceived risk
- The ability to develop rotational plans for periodic reviews, based on the rated risk levels

Figure 6.3—Risk Heat Map

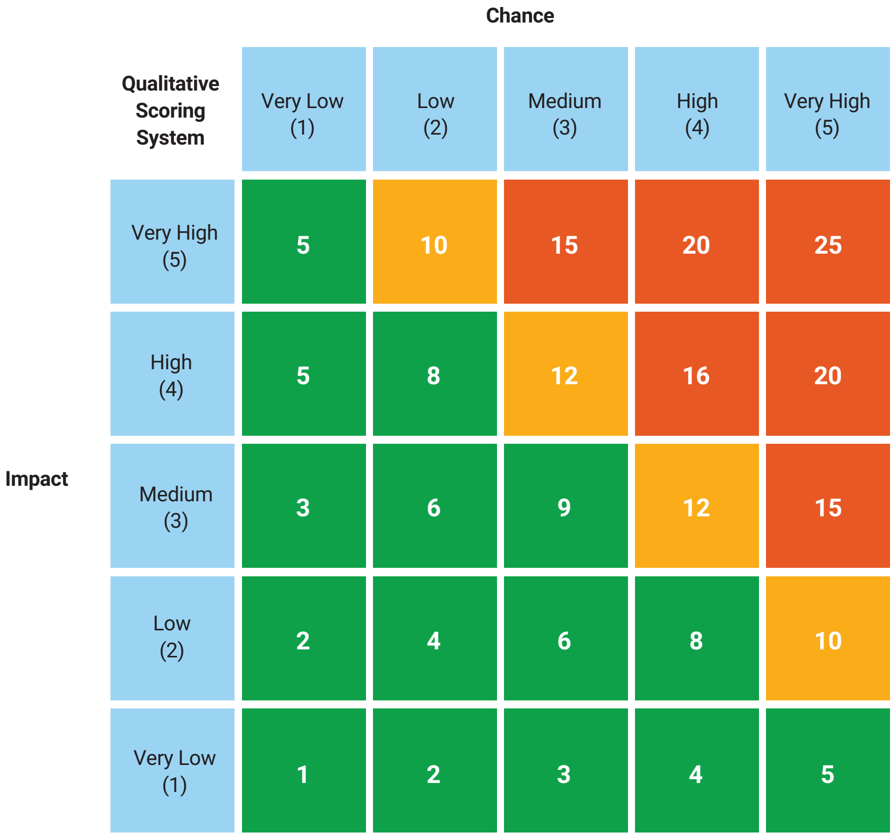

		Chance				
	Qualitative Scoring System	Very Low (1)	Low (2)	Medium (3)	High (4)	Very High (5)
Impact	Very High (5)	5	10	15	20	25
	High (4)	5	8	12	16	20
	Medium (3)	3	6	9	12	15
	Low (2)	2	4	6	8	10
	Very Low (1)	1	2	3	4	5

This approach allows each organizational unit to maintain a portfolio view of the risks associated with the cloud services it is responsible for, and develop a tailored assurance plan.

6.6.4 Evaluate Compliance Program Reporting on Risk and Controls to Various Stakeholders

The compliance program needs to provide feedback to the stakeholders on the results of the risk evaluation in the form of a report identifying the stakeholder risk.

- Example stakeholders follow:
  - C-suite officers (CISO, CIO)
  - Application owners
  - Customer
- Example compliance report topics follow:
  - Missing controls
  - Recommendations for improvements
  - Status of key risk indicators

The company needs to identify organizational ownership and responsibilities for the lifecycle of reporting and tracking risk. This means setting up a feedback loop, proper communication paths, and internal compliance checks. The internal compliance team should be checking with the control owners regularly and providing feedback if they are not compliant. All of this provides assurance—helping management understand how the reports are coming in and what their risk posture is to inform their decision-making.

There should be two-way communication with stakeholders and control owners. Having a strong cadence of communication is essential to reduce and account for potential risks in the future. If one party can identify a potential threat or has heard about the latest vulnerabilities, having strong communication channels to inform the other party will be beneficial for both sides. With technology changing rapidly, joint efforts are essential to success and risk reduction now and into the future.

An existing risk management process needs to be tied to the cloud infrastructure in use. Education of risk teams on cloud security is recommended—and not only because regulators expect organizations to have knowledgeable people on their teams. The cloud auditor checks how the existing risk processes are dealing with cloud-specific risk, such as the frequency of occurrence or what happens over time. A server that was not properly patched last week might no longer exist because of a redeployment. Is the organization's risk management process ready for this fast-paced change?

Consider establishing a strategic information hub function comprised of a group of people across multiple teams that monitor the cloud for compliance, providing oversight across all cloud environments and across all DevOps teams. Some aspects to consider in connection with this monitoring: cost, ITIL processes, DevOps maturity and security compliance.

For example, it is not only important to discover the existence of unpatched servers, but also interesting to analyze how many days the servers are not patched and investigate the reasons. The cloud auditor should be alert to issues like wrong priority setting of a product owner or structurally spending all team capacity on change and not on operational security tasks or basic hygiene. Of course, this depends on the scope and depth of the audit.

Monitoring and Keeping Under Control the Risk of Noncompliance

Capturing all risk in a central risk register will help address concerns with noncompliance issues. It allows an organization to keep an eye on risk it has accepted, and revalidate them after a certain threshold is met. Depending on the risk level, that threshold may vary. Risk that is accepted should also have compensating controls in place to help mitigate risk exposure. The cloud auditor should verify if the risk can be a business-as-usual (BAU) process with all security settings applied.

Central risk register solutions will need to align with compliance and governance programs to help constantly update changes or modifications. Consider onboarding your cloud risks in one of the available GRC suites your organization probably has available. Without having the governance, risk and compliance programs align with each other, there is no good way to measure for noncompliance and track open items. Having outdated and nonaligned programs will introduce more risk to the company and can do more harm than good. It is important to know the current compliance status at all times to keep risk under control.

Evaluate How the Program Organizes Multiple Levels of Risk Management and Skills Staffing

The company needs to have different teams in place to deal with risk issues. Identify which people will be the ones to remediate or accept, and which people will be able to talk through the core of the problems—e.g., technical SMEs versus someone who could explain the issue to management in a less technical and more accessible manner.

First-line workers will generally be those who are using the control solutions, e.g., the SME on the servers who realizes the machine has been altered or modified. They will follow the proper escalation and may even need to report the issue as an incident. The second-line workers will generally be those who receive the ticket from the first-line SME. These may include the SOC/CSIRT teams. They will begin to review their monitoring consoles for any new alerts or suspicious activity found in the environment.

Together, these two lines of defense will collaborate to mitigate the issue and overall reduce its risk of impact. In a DevSecOps way of working, the DevOps team is expected to partially solve security incidents or noncompliance issues. DevOps team members collaborate with their colleagues from the SOC, who usually notice security noncompliance first.

The order of operations for these tiers is important only from the hierarchy perspective. When something goes wrong, whoever discovers the issue will end up reaching out to the SMEs on the systems to help fix the issue.

Managers will get involved for escalation of incident handling and larger decision-making. For smaller incidents, decisions can be made right away by the teams that are hands-on to the incident to help move the case forward. This could include actions by both the SMEs and the SOC/CSIRT team. Depending on the severity of the issue, decision-making may be escalated to the most senior executives, who are generally looking for updated status reports on the issue and how and when it will be resolved. They will want to know the severity and impact of the issue and aim to solve it with as little damage as possible.

6.7 Introducing Services Changes (New Services, Markets, etc.)

The service architecture will change over time, either because new services are introduced or because existing services change. Regardless of the source of change, the organization needs to make sure that any evolution of the service portfolio and architecture does not have a negative impact on its compliance posture.

6.7.1 Evaluate Threat and Risk Modelling for Development and Operations

Besides considerations about evaluating CSPs, one of the most important evaluations concerns service development and operations, and the threat and risk modeling connected to them.

As part of evaluating the compliance program, the organization will need to ensure that development and operations are following a mature threat and risk model. This applies to both development and operations within the organization and at the CSP. In general, it is advisable that the assessor verify that the organization follows the security-by-design principle, and that security is embedded into the product or service design and development process, as opposed to being added post-development or as an afterthought.

Assessors should consider the development methodologies being used to build an application, which will typically be either waterfall or agile/DevOps.[254] Both approaches are effective means of producing software and services in a secure way. The assessor needs to be aware of how both models work on paper and in practice, and the key points along the route they need to audit and check. An assessor also needs to consider that some organizations will use both models. For example, an organization may have old products built using one development methodology with the associated controls and governance, and undertake new projects using the other approach.

Combining the concept of security-by-design with the agile operating model will help an organization develop and meet the intended goals associated with its intended risk model or maturity framework. In development, the secure development life cycle (SDLC) is a process for planning, creating, testing and deploying an information system, as shown in **figure 6.4**. As an organization prepares to introduce new solutions, the SDLC will provide a structured approach for the key steps to take.

[254] See chapter 8.

Figure 6.4—Software Development Life Cycle (SDLC)

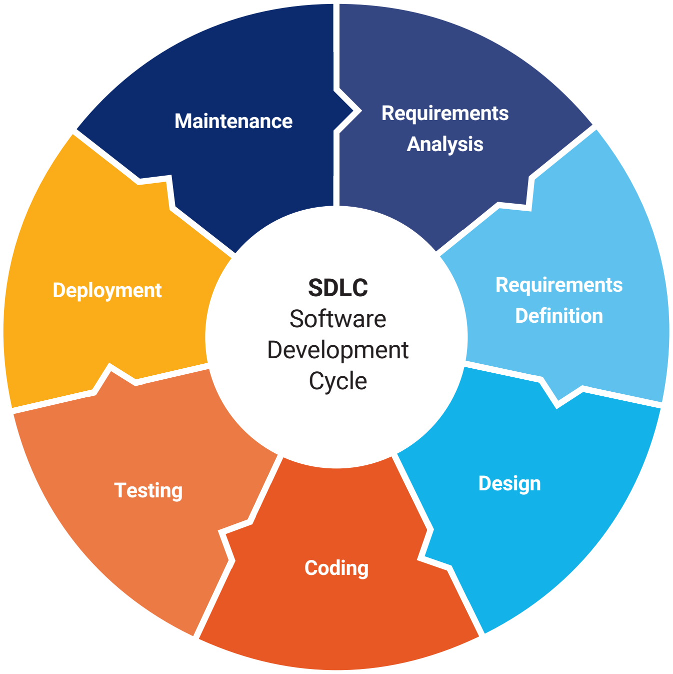

Having compliance requirements ahead of time will allow the teams to embed them directly into the development instead of after the project has been completed. Adding security later is typically more expensive and less effective, to the extent that some requirements might not ever be satisfied.

Tip: A great example of designing security into the development process is sharing API keys. Having an API that enables push-and-pull capabilities to extract or enter data from the application is extremely popular and beneficial. Knowing that API keys will need to be stored and shared once the application is in production will allow the teams to embed a secure storage and sharing methodology into the design of the product while building it.

A customer that needs to add a feature to an existing product may work with the provider to include a bolt-on solution as an accommodation for the intended requirement, e.g., providing BYOK capabilities. The provider needs to be upfront about any potential issues or flaws that the bolt-on solution may have. It is possible the provider is not aware of any risks. The customer needs to test and validate that the bolt-on solution works as intended, and evaluate whether an improvement model will be provided if there are issues. As a customer, it is very important to ensure that adding new features does not accidentally cause a greater amount of risk to the environment. To be on the safe side, documenting these issues in the contract or SLA helps place the responsibilities on the intended party.

The cloud also introduces the concept of Infrastructure as Code (IaC) i.e., a practitioner defines the cloud services and their configuration settings in a source code file. It is then possible to create disposable cloud environments by using a cloud orchestration tool that takes the IaC blueprint and turns it into a cloud environment, containing everything required for that service or application to run. From the orchestration tool, assessors can gain visibility into build scripts and see implementation decisions around the services and code being used.

6.7.2 Shared Responsibility Implications

This section invokes the shared responsibility model. Where there is an application that consumes multiple third-party cloud services, the assessor's goal is to establish whether the organization has identified security controls of the cloud provider, investigated relevance of controls for developing secure software, and obtained evidence to support the answers.

Questions to ask include the following:

- Does the CSP follow secure coding standards? Does it train its developers to do so? Do those standards apply to the language the application is built in?
- Does the CSP perform threat modeling?
- What end-to-end security tests does the CSP carry out?
- What static or dynamic code analysis tools are used to identify potential vulnerabilities quickly?
- How does the CSP conduct peer reviews for security-sensitive code modules?
- What third-party security libraries does the CSP use, and does it track how it keeps them up to date?
- What penetration tests does the CSP perform on application features to be deployed before putting them into production?
- How does the CSP confirm that developers do not have access to production data?
- How does the CSP confirm that support staff must request access to customer data in advance?
- Is there an external vulnerability reporting mechanism for security researchers and a published policy on how the CSP will treat security researchers who approach it with vulnerabilities?
- What change control process does the CSP use to assess the potential security impact of the change and to engage the right security experts to respond and contribute before a change is made in production?
- Do the CSP's developers keep secrets outside of the codebase and in a secure repository?
- Is access to the codebase limited on a need-to-know basis?
- How does the CSP detect malicious code planted in a source code repository before it is deployed in production?

6.7.3 Compliance Tolerance to Service Changes

An enterprise can only outsource so much of its risk tolerance to a third-party service provider. When the provider begins to tailor its services without strict adherence to the contract, or when it introduces new changes to its products and solutions, the customer needs to understand if the changes (either technical or contractual) create any disruption in the governance, risk management and compliance posture of the organization. Customers need to continuously monitor the providers' business models for change.

The customer needs to decide for itself what levels of provider compliance tolerance and service level tolerance are acceptable before making any changes. Once those thresholds are met, the customer needs to adapt and begin evaluating additional service providers that offer the value the company needs for its current business model, not yesterday's.

Example: Consider the case of a service provider that collects users' mobile phone numbers for the purpose of offering two-factor authentication access. the organization uses this cloud service and is happy with the added security provided. The service provider later decides to use this information for targeted advertising,[255] allowing users to opt out. Even with the opt-out choice, this change may be illegal in some jurisdictions and may also create friction with the customers of the organization who are subject to the change. Even though the organization likes the product, it is forced to adapt and will need to begin evaluating additional or other providers that meet its existing requirements. Keeping the current provider would introduce more risk to the business.

[255] Kastrenakes, J.; "Twitter faces $250 million FTC fine for misusing emails and phone numbers," The Verge, 3 August 2020, www.theverge.com/2020/8/3/21353232/twitter-ftc-fine-misused-email-phone-advertising-2011-settlement

The only way to be prepared for these types of situations is to have connected risk, compliance and governance programs. As these three programs align, any modification to one component will trigger alerts in the others. Some providers maintain up-to-date information on the latest industry requirements relevant to their solutions. It may be hard to track all these changes in-house. Having multiple sources of communication will increase the likelihood the enterprise will be notified as soon as possible.

6.8 Evaluate the Need for Continuous Assurance/Continuous Compliance

Traditionally, evaluating compliance has been a point-in-time or period-of-time exercise, regardless of whether it was an internal or an external audit. The evaluation was performed at a certain given time and would produce a snapshot of the status quo of a certain process or service. Some typical examples of such an approach are an ISO/IEC 27001 certification process or a SOC 2 attestation engagement. Both provide a detailed picture of an organization's security and compliance posture (limited to the scope of the audit, of course) that can be used to make assumptions about its near future posture.

None of the static approaches to compliance can give assurance of what could happen one day after the audit is finalized. Cloud services change rapidly—sometimes even several times in a day—and depending on the level of assurance the cloud customer requests, a different approach might be required, based on continuous auditing.

6.8.1 Continuous Auditing and Continuous Monitoring

The concept of continuous auditing is often confused with continuous monitoring, so it is important to clarify the differences. CSA describes continuous auditing as an ongoing assessment process that aims to determine the fulfillment of service qualitative objectives (SQOs) and service level objectives (SLOs), conducted at a frequency requested by the purpose of audit.

In the field of information security management, continuous monitoring is defined as maintaining ongoing awareness of information security, vulnerabilities and threats to support organizational risk management decisions.[256]

Continuous auditing is the evaluation of control implementation by auditors, whereas continuous monitoring describes the continuous feedback provided to management regarding key processes. These two concepts overlap, but they target different stakeholders in an organization.

The results are typically expressed as a percentage, or a range. Examples could be evaluation of the service availability, or of Recovery point objective RPO[257] or RTO.[258]

Control objectives also can be translated into qualitative measures. In this case, the auditing will be manual and the results expressed as a nominal scale (yes/no/not applicable) or as an ordinal scale (high, medium, low). Examples could be the evaluation of a security policy.

In general, the fact that a certain control can be expressed as SQO or SLO will impact the way the control can be audited (manual vs. automated) and ultimately the possible frequency of the checks.

[256] Dempsey, K.; N. Chawla; A. Johnson; R. Johnston; A. Jones; A. Orebaugh; M. Scholl; K. Stine; "Information Security Continuous Monitoring (ISCM) for Federal Information Systems and Organizations," National Institute of Standards and Technology, September 2011, https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-137.pdf

[257] Recovery point objective

[258] Recovery time objective

Of course, the frequency is impacted not just by possible technical limitations but also by the risk appetite, the compliance requirements, and the budget of the enterprise that sets the policy.

6.8.2 Continuous Auditing and Continuous Assurance

Continuous assurance in the cloud is often associated with services provided to customers operating in highly regulated or critical business sectors, such as finance and healthcare.

The concepts of continuous assurance and continuous compliance are closely related. In essence, continuous assurance is meant to be the result of continuous monitoring and auditing. The simple logic behind it is that reducing the time interval between one evaluation/test/audit and the next, reduces the risk that something could go wrong (e.g., a new vulnerability, a misconfiguration) and therefore increases the amount of assurance that the service is performing as promised.

The interval between checks typically depends on the nature of the security control/measures to be evaluated. That is why CSA relates the idea of continuous auditing to the concepts of SLOs and SQOs.

As shown in **figure 6.5**:

- Evidence is collected from the information system.
- A measurement applied to that evidence according to a metric produces results.
- The measurement results are compared to the SLOs or SQOs to determine whether the objective has been met.

Continuous auditing describes the process of collecting evidence and assessing it to obtain a quantitative (SLO) or qualitative (SQO) result for the purpose of evaluating a system's performance.

Figure 6.5—Measuring Service Level Objectives (SLOs)

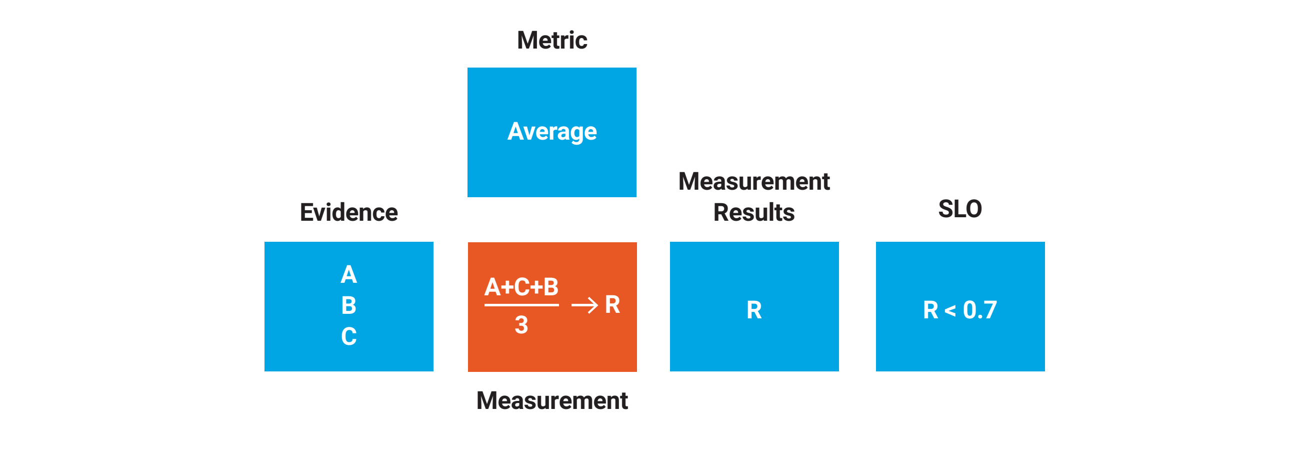

The process described plays out in different ways and produces a different auditing approach and different results depending on the nature of the security control objective used as input. In the case of a control objective that can be translated in a quantitative measure, the auditing can leverage automated tools and happen automatically, and the results are typically expressed as a percentage or a range. Examples are evaluation of the service availability or of recovery point objective (RPO) or recovery time objective (RTO), etc.

The other scenario is one of control objectives that can be translated into qualitative measures. In this case, the auditing will be manual, and the results will be expressed as a nominal scale (yes/no/not applicable) or as an ordinal scale (high, medium, low). Examples are the evaluation of a security policy.

In general, the fact that a certain control can be expressed as SQO or SLO impacts the way in which such a control can be audited (manual vs automated) and ultimately the possible frequency of the checks.

Of course, the frequency is not just impacted by possible technical limitations, but also by the risk appetite, the compliance requirements and the budget of the organization that sets the policy on how often a certain control shall be evaluated to provide the necessary level of assurance.

Depending on the level of assurance the organization would need to obtain from its CSP and the level it would need to provide to its shareholders, customer and regulators, the assessors should focus on understanding the following:

- Whether the frequency of the control evaluation is aligned with the level of assurance required
- Whether the controls are designed to meet the requested testing and evaluation frequency (e.g., if the organization follows a DevOps approach, are the controls natively designed into the software?)
- Whether the source of evidence allows for continuous auditing (Are logs and feeds available with the level of details and frequency required? Are the logging and monitoring capabilities natively provided by the CSP sufficient, or do external tools need to be integrated? Is the reporting from the CSP, e.g., incident reporting, aligned with the organization's compliance needs?)

Even with continuous assurance and compliance there may be configuration drift.

Measuring Drift and Configuration Changes

Continuous compliance allows the company to make sure that controls are in place and that systems are not being put in place without the correct controls. Change management and real-time reporting need to be enabled, to provide notification when drift occurs. This drives assurance that the company remains compliant with its controls in place. There are tools that can help organizations perform these functions, such as Salt Stack, Ansible and Chef/Puppet. Ultimately, it is the real-time remediation or risk acceptance that drives assurance and peace of mind for the company.

Example: When a company performs configuration management on its endpoints or one of its servers, what happens when a developer pushes something out to production? The developer may not have realized that one of its security configurations became modified. Does the company have a way to enable continuous compliance checking and validation to confirm the control is still in place? Otherwise, it may drift away from a compliant state. Does the company have real-time remediation and change management?

Page intentionally left blank

6.9 Chapter 6 Knowledge Check

REVIEW QUESTIONS

1. What is the **MOST** important reason to use an external auditor?

   A. To identify control gaps, taking an independent/external point of view from the organization
   B. To avoid conflicts within the organization, which can arise when an internal auditor is perceived to criticize the work of colleagues
   C. To reduce costs in small organizations by eliminating internal auditors
   D. To provide guidance when implementing DevSecOps in an organization that uses the cloud extensively

2. A shadow cloud service typically:

   A. Avoids or evades appropriate risk management
   B. Evades due-diligence assessments
   C. Creates weaknesses in the organization's compliance posture
   D. Provides compliance benefits

3. Solutions within an organization should be aligned to the requirements it needs to meet. What is the **BEST** way for a customer to monitor this?

   A. The customer should have a program and process in place to validate it, at least on a yearly basis.
   B. Onsite audit once a year
   C. Fill out a CAIQ and submit to the STAR Registry
   D. Third-party certification

4. How does the company prepare their compliance program for instances they cannot account for (i.e., breaking news or published articles)?

   A. Have strong understanding of their customer base
   B. Have strong understanding of all their sources of risk to the organization
   C. Have strong ability to communicate to vendors
   D. Have inside knowledge of what their competition is working on

5. What type of changes should be monitored by customers using cloud services, when considering compliance and risk? (Select all that apply.)

   A. Monitor changes introduced by the service provider, which present new compliance risk for the customer
   B. Monitor new requirements introduced by the customer, which present compliance risk due to misalignments with the service offered by the cloud provider
   C. Monitor changes introduced by the service provider that are also offered by cheaper competing cloud service providers
   D. Monitor requirements introduced by the customer that are different from the requirements introduced by the cloud provider

6. What is continuous auditing?

   A. An ongoing assessment process that aims to determine the fulfilment of service qualitative objectives (SQOs) and service level objectives (SLOs), conducted at a frequency requested by the purpose of audit
   B. Maintaining ongoing awareness of information security, vulnerabilities and threats to support organizational risk management decisions
   C. The continuous feedback provided to management regarding key processes
   D. The amount of assurance that the service is performing as promised

Answers on page 324

Chapter 6 ANSWER KEY

Correct answers appear in **bold** font.

1. **A. Correct, an external view can reveal issues that are sometimes invisible to an organization.**

   B. Incorrect, professionalism is the key to avoid conflicts.

   C. Incorrect, external auditors are not designed to replace internal auditors.

   D. Incorrect, DevSecOps does not specifically require an external auditor.

2. **A. The shadow cloud service typically avoids or evades appropriate risk management and/or due diligence assessments.**

   B. The shadow cloud service typically avoids or evades appropriate risk management and/or due diligence assessments.

   C. The shadow cloud service typically creates weakness in the organization's compliance posture.

   D. The shadow cloud service typically creates weakness in the organization's compliance posture since formal policies are not applied.

3. **A. The customer should have a program and process in place to validate, at least on a yearly basis, that the solutions within the organization are aligned with the requirements it needs to meet. The customer will need**

3. **A. The customer should have a program and process in place to validate, at least on a yearly basis, that the solutions within the organization are aligned with the requirements it needs to meet. The customer will need to have a vendor management program, an approved software list explaining the business need, and proof the vendor is complying with the requirements.**

   B. A provider may provide a white paper that explains specifics about controls available. The white paper may not go into enough detail and just highlight key focus areas as a marketing approach to draw the customer in. The customer will need to contact the provider and get more specific information to see if the CSP is able to meet the specific business need that is intended.

   C. The CIAQ is for CSPs, not customers.

   D. Third-party certification of a customers process may validate their process but it will not validate the CSPs process.

4. A. Knowing the customer base for the organization will help in configuring a compliance program but not help prepare for incidents such as breaking news or published articles.

   B. Knowing what the organization is responsible for and maintaining current statistics on what are the risk to the industry, business, etc., will help the compliance program know what controls are needed to be in place to protect and prevent risk exposure.

   C. Being able to communicate in general is a very strong trait to have, although when instances occur, such as breaking news or published articles, speaking with vendors is not going to be relevant. If anything, the company would need to speak to the public to address any issues.

   D. Understanding how the competition handles their breaking news or any published articles would be very helpful, although knowing what they are working on is not going to help in this case.

5. **A. Correct, these changes need to be scrutinized by the customer.**

   B. Correct, the customer needs may change and the features offered by the cloud provider may not match the customer's requirement anymore.

   C. No, that's unrelated to security.

   D. No, requirements are matched with offerings.

6. **A. Continuous auditing is defined as "an on-going assessment process that aims to determine the fulfilment of Service Qualitative Objectives (SQOs) and Service Level Objectives (SLOs), conducted at a frequency requested by the purpose of audit".**

   B. This is the definition of continuous monitoring.

   C. This is a description of continuous monitoring.

   D. This is a description of continuous assurance.

Chapter 7:

CCM Auditing Guidelines

CCM Auditing Guidelines

CCM Auditing Guidelines

7.1 Learning Objectives

After completing this chapter, learners will be able to:

1. Detail CCM Auditing Guidelines.
2. Define the CCM Audit Scoping Guide.
3. Explain the approach taken in the CCM Risk Evaluation Guide.
4. Evaluate the CCM Audit Workbook.
5. Apply the CCM Auditing Guide.

7.2 Overview

The CCM Auditing Guidelines provide a baseline understanding of the Cloud Controls Matrix (CCM) audit areas and provide tools and resources to auditors performing a CCM-related audit and assessment. The information in this chapter represents a generic guide, and the auditor will need to customize the descriptions, procedures, risks, controls and documentation to address the specific objectives of the audit.

These guidelines are intended to support cloud security audits or assessments (either internal or external) of organizations of any size, business, cloud deployment complexity and maturity. The target audience is the same as the audience of the CCM. This audience includes auditors who plan to perform audits against CCM on cloud service providers (CSPs), cloud customers using the CCM framework to evaluate their cloud service portfolio, and CSPs or cloud customers that intend to use the CCM framework to guide the design, development and implementation of their cloud security controls.

The CCM Auditing Guidelines include the following sections:

- **CCM Audit Scoping Guide**—During the audit planning stage, this section can be used to provide guidance on how to define the scope of an audit.
- **Risk Evaluation Guide**—This section is a framework to assess the risk level within an organization.
- **CCM Audit Workbook**—This section offers a workbook to facilitate and guide a CCM audit.
- **Request List**—This section provides a list of items that could be requested by an internal or external auditor to support a CCM audit.
- **Exception Report**—This section provides a template to document exceptions identified during a CCM audit.

7.3 CCM Audit Scoping Guide

During the audit planning stage, cloud auditors can use the CCM audit scoping guide to help define the scope of an audit. Scoping an audit includes identifying parameters that are relevant for building a sample and the topics to address. This scoping guide provides key factors to consider but is not intended to cover all scoping criteria. None of the factors should be considered mandatory. The auditor will need to customize the criteria to address the objectives of the specific organization audit.

7.3.1 Scoping Key Factors

There are several key factors to consider when determining an audit scope. The first key consideration is what is being audited? Who is being audited—the CSP or the customer? Who is performing the audit? The next consideration is the purpose of the audit. Why is the audit taking place? Is it a third-party audit related to a STAR Certification or Attestation? Is it a first- or second-party audit to evaluate the compliance posture against the CCM requirements requested by the board, management or auditing committee? Is it an exercise driven by internal audit as a part of the overall organization audit plan?

Once the purpose of the audit is clear, the auditor determines the audit scope by asking the following questions:

- What type(s) of cloud services have been implemented (IaaS, PaaS, SaaS) and which ones are in the scope of the audit?
- What are the (critical) assets for which the auditor seeks assurance?
- What is the deployment model (private, public, hybrid), and in the case of a private cloud, who is managing the infrastructure?
- How mature is the auditee organization overall from a cloud governance and security perspective?
- What is the nature of the applications in (or being moved to) the cloud? What additional aspects need to be considered?
- What data will be processed in the cloud?
- What parts of the organization are in scope (e.g., head office, core IT teams, subsidiary companies)?
- What teams or perspectives are in scope (e.g., teams for core cloud infrastructure management, DevOps, architecture, security)?
- Has the organization considered the SLA or other contractual requirements?

This list of basic questions can be extended with additional ones found in chapter 2 and with the guidance on scoping included in chapter 6.

Cloud Service Model

Each cloud service model uniquely impacts the scope of an audit. Following are the three main cloud service models that organizations encounter, and the factors to consider when scoping an audit for each. These factors should be integrated with the additional considerations provided in chapters 2, 3 and 5.

With Software as a Service (SaaS), the organization has the least amount of control and the least amount of direct responsibility for the application served by the CSP. However, even though the organization is not directly responsible for the bulk of the controls, it is still accountable and may be responsible for things such as access controls. In addition, most organizations have their own internal controls or processes for running and maintaining the SaaS service.

With Platform as a Service (PaaS), the CSP serves the platform upon which the organization can develop and deploy new applications. Although the CSP does manage certain controls related to the SDLC process, the bulk of the responsibility for developing software that is both secure and operationally sound falls on the cloud customer. Most enterprises have their own cloud service infrastructure team that makes the PaaS services available to the DevOps teams in a secure and compliant fashion. Some organizations outsource this infrastructure maintenance to a managed cloud service provider.

With Infrastructure as a Service (IaaS), the organization is responsible for most IT operations and information security processes. Although the CSP serves the core infrastructure, it is the organization's responsibility to maintain computing resources in an operationally sound and secure manner.

When considering the scope of the audit, auditors should take the following considerations in account for each of the three cloud service models listed previously:

- Alignment with the business and IT strategy

- Governance, risk assessment and compliance
- Shared responsibility model, accountability and boundaries
- Third-party management
- Cybersecurity operations
- Cloud data protection and lifecycle management

Additional details are provided in chapter 2.

Special considerations for SaaS audits are:

- Third-party management, and in particular the contracting and vendor monitoring processes (see chapter 5)
- IAM in the cybersecurity operations (see chapter 3)
- Accurate SaaS configuration
- The shared responsibility model

Cloud Deployment Model

The next step in assessing the scope of the cloud audit is determining where the cloud services are hosted. Following are the different cloud deployment models and how each model will impact an audit scope.

With a private cloud, the organization controls the entire infrastructure, including hardware, software, facilities, administrative personnel, security controls, etc. Therefore, the organization inherits all the risks from a traditional on-premises environment. The audit could focus on assessing all applicable security aspects, including but not limited to personnel threats, natural disasters, external attacks, regulatory noncompliance and malware.

With a public cloud, the organization loses control over infrastructure, oversight and enforcement. The audit should focus on the components that the organization can manage directly, on the controls for which the responsibility is shared with the CSPs, and on those components and the capabilities on which the organization needs to exercise a duty of care, such as identity and access management, data protection, or incident management or BC/DR planning, etc.

A hybrid cloud inherits the characteristics of a public cloud and private cloud. The audit scoping should factor in considerations that apply to both.

The multi-cloud model uses multiple public cloud providers to deliver the same type of service. The organization should evaluate all considerations noted for each individual cloud provider.

A hybrid multi-cloud uses multiple public cloud providers to deliver the same type of service in addition to using private cloud and traditional on-premises IT. The organization should evaluate all considerations noted for each individual cloud provider.

Maturity of the Cloud Environment

The audit needs will be defined by the maturity of the organization adoption of cloud services. An organization that is new to the use of cloud will have different requirements and areas of focus than one that has been using cloud services for some time.

Early Cloud Adoption

If the organization is moving its first application or set of applications to the cloud, the auditor should first evaluate the organization established baseline security capabilities, tool sets and metrics in aligning with cloud adoption security risks. It is likely that the organization has no information security governance or processes in place specific to the cloud at this time, so a traditional audit will not provide much value, nor will it address the most pressing risks. Therefore, the auditor should tailor the audit to focus on the security risks associated with cloud adoption and consider audit procedures designed to provide assurance over the broader set of cloud risks at the enterprise level.

Cloud Implemented but With Immature Processes

In this scenario, the organization has already implemented one or more applications in the cloud, but it has done so in an ad hoc fashion without centralized governance. In this circumstance, the auditor should first assess the existing cloud services for any significant information security-related gaps that could impact the business. Second, the auditor should look at audit procedures that assess security- and governance-related controls.

Cloud Implemented and Mature

In this scenario, the organization is using cloud computing technology extensively as part of its business strategy. It should already have formal processes and frameworks (such as the NIST Cybersecurity Framework) to assess third-party services and manage the internal security control framework. Depending on the request, the auditor may want to align the CCM controls with the existing controls the organization has defined.

The auditee has already implemented the CCM controls (or the subset of them relevant to the scope and required, given the risk exposure). The auditor has to determine if controls the auditee has selected are relevant and correctly implemented.

Intended Use of Cloud Computing and the Inherent Characteristics

Auditors can use this section of the guide to help identify the nature of the applications in the cloud or being moved to the cloud, and any additional aspects that may need to be considered for the audit.

- **Security Considerations**—Focus on inherent privacy and security risk coming with the current or intended use of cloud. For example, is the data highly sensitive? Are the controls around data protection sufficient? Are there concerns over segregation of duties? Are the appropriate access controls in place?
- **Compliance Considerations**—Focus on inherent compliance risk associated with the current or intended use of cloud. For example, are the systems or data impacted by regulations such as GDPR, SOX, GLBA, HIPAA, ITAR, etc.?
- **Reliability Considerations**—Focus on inherent reliability risk associated with the current or intended use of cloud. For example, can the CSP ensure systems, data and functionality are available when needed and the required level of performance is maintained? Are the cloud services appropriately resilient to disasters? Does the CSP BC/DR plan meet the customer SLA requirements?

7.4 CCM Risk Evaluation Guide

This section provides a framework that an auditor can use to assess the risk level within an organization.

7.4.1 CCM Risk Evaluation Approach

The following three-step approach to assess the organization risk level is a simplification of the processes described in chapters 1 and 2.

1. Establish the organization cloud risk profile with risk-evaluation questions (**figure 7.1**).
2. Identify and document the key risk categories for the organization in line with the organization line of business and geography, its use of cloud, and the unique requirements to which it has to adhere (**figure 7.2**).
3. Evaluate and document this risk in the section provided in the CCM Audit Guidelines (**figure 7.3**).

Figure 7.1—Step 1: Establish Cloud Risk Profile

Establish the organization cloud risk profile with risk evaluation questions. (Following is a list of sample questions.)
Example Starter Questions How is your cloud technology aligned with the overall business strategy? Has the information security policy been updated to reflect could technology management?

Figure 7.2—Step 2: Identify Key Cloud Risk

Identify key risk categories that the organization has to face and then document risk within each category, in line with the organization line of business, its use of cloud and unique requirements to which it had to adhere.

Example Risk Category	Example Risk
Strategy, Business Case and Business Requirements	Cloud strategy is not clear in the organization and does not have enough support from the organization governance bodies (tone from the top)

Figure 7.3—Step 3: Evaluate and Document Risks

After the organization risk profile is built and risk is identified, each risk needs to be individually analyzed for:

1. **Potential Impact**—The organizational impact (in the form of financial loss, reputation damage, etc.) if an event of the risk occurs
2. **Likelihood**—The probability of occurrence of an impact that affects the organization
3. **Risk Mitigation**—Mitigation put in place to reduce the likelihood of risk and potential impact
4. **Residual Risk**—Any risk left after mitigations are in place

Finally, an overall risk level should be determined based on the analysis performed

Impact of Risk	Likelihood of Risk	Risk Mitigation

After the organization risk profile has been built and the risk identified, each risk condition needs to be individually analyzed to determine the potential impact of the risk condition, the likelihood of the risk condition happening, what should be done to mitigate the risk condition, and residual risk. Finally, the overall risk level should be determined based on the analysis performed. The auditor is responsible for customizing the risk-evaluation framework to address the objectives of the audit. See chapter 1 for recommendations on a general risk-evaluation process and best practices to use.

7.4.2 Establishing Cloud Risk Profile

The first step in risk evaluation is to establish the organization cloud risk profile with the risk evaluation questions. Consider the questions in **figure 7.4** as a starting point for organizations to evaluate their risk profile. These are intended for auditors to expand on or adjust to better fit their organization.

Figure 7.4—Example Starter Questions
How is your cloud computing technology aligned with the overall business strategy?
Has your information security policy been updated to reflect cloud computing technology management?
Does your third-party risk assessment include specific risks associated with cloud computing technology (e.g., shared roles and responsibilities, compliance requirements, right to audit)?
What legal/regulatory requirements apply to your organization's use of cloud computing technology (e.g., based upon data privacy and security laws)?
Have you completed a cloud migration risk assessment, and were cloud computing technology-specific risks considered?
Do you have an inventory of all CSPs and associated contracts, and third-party assessment and attestation reports available?
Has your organizational BC/DR plan been updated to reflect the cloud adoption?
Has your organizational privacy policy been updated to reflect the situation post-cloud migration?
Has the organizational incident management policy been updated to reflect the situation post-cloud migration?
Has the organizational secure software development lifecycle (SSDLC) been updated to reflect the cloud migration?

7.4.3 Identify Key Risk Categories

The next step is to document the organization key risk categories and the risk conditions within each category based on the organization line of business, how it uses the cloud, and any unique requirements.

Figure 7.5 provides a list of common cloud risk factors for each risk category included in the guide. The auditor can use these examples as a starting point to identify the risk and risk categories that the organization should address.

Figure 7.5—Common Cloud Risk Factors

Example Risk Category	Example Risk Description
Strategy, business case and business requirements	Cloud strategy is not clearly embedded in the organization and does not have enough support from the organization's governance bodies (tone from the top).
Governance function	The governance function to provide meaningful and sustainable process and oversight over cloud technology is not established.
	The organization's asset inventory does not include cloud technology assets and does not document them to ensure governance function coverage over a complete list of assets.
	The enterprise-wide awareness and training program does not include cloud knowledge management, and skills do not exist to effectively build and manage cloud technology.
	Business continuity and disaster recovery is not considered.
Risk management	The enterprise risk management function does not include cloud-specific considerations in its assessment of enterprise-wide risk and impact.
	The cloud data management function is not integrated into enterprise-wide risk assessment to ensure information security and privacy goals are achieved.
Compliance	The organization is not using a unified framework to integrate cloud compliance with regulatory requirements.
	Contractual obligations to address cloud compliance requirements are not documented, approved and monitored.
	Certifications with global security standards specific to the cloud are not reviewed, and the impact of noted findings is not assessed.
Third-party management	Outsourced cloud service experiences interruption, breach or loss of data stored at the CSP.

Figure 7.5—Common Cloud Risk Factors *(cont.)*

Example Risk Category	Example Risk Description
Cloud data protection and life cycle management	Cloud data inventory does not exist or is not updated or classified to ensure appropriate handling.
	Cloud service and data lifecycle management does not consider security risk.
	No policies and solutions are used for data loss prevention in the cloud environment and during transit.
Identity and access management	Access to cloud technology is not considered in corporate access management policies and procedures.
	Access to cloud technology is not appropriately restricted.
Application security and development operations	Application architecture and configurations do not incorporate cloud security measures.
Portability and interoperability	New cloud technologies and services are not integrated, and data is not portable between legacy systems and existing cloud solutions.
Incident response, notification and remediation	Cloud operations and security incidents are not monitored and properly resolved.
Threat and vulnerability management	Vulnerabilities in the cloud environment are not identified and remediated.
Key management	Encryption is not established to protect cloud data from unauthorized access.
Resilience (backup and recovery)	Data in the cloud environment cannot be recovered in the event of disaster.

7.4.4 Evaluate and Document Risk

The last step in this section of the guide is to evaluate and document the risk. Once the organization risk profile is built and the risk events are identified, each risk event needs to be individually analyzed for the following:

- **Potential impact**—The organizational impact (in the form of financial loss, reputation damage, etc.) if a risk event occurs
- **Likelihood**—The probability of an occurrence that affects the organization
- **Risk mitigation**—Mitigations that should be put in place to reduce the likelihood of risk and potential impact
- **Residual risk**—Any risk that will be left after the mitigations are put in place

All of this should be documented in the audit guideline spreadsheet, and then the overall risk level should be determined based on the analysis.

7.5 CCM Audit Workbook

The CCM Audit Workbook is intended to facilitate and guide a CCM audit. Auditors are responsible for tailoring the audit work programs for the organization and should validate all the included processes and control information with the organization before executing any work.

For each CCM control, the workbook contains a place to document the following:

- Control audit frequency
- Test plan
- Test result
- Reference to supporting documentation

- Identity of who performed the audit
- Control audit conclusion

7.5.1 CCM Control Specification

This step describes the purpose of the CCM control and which issues or challenges it addresses within a particular organization. See **figure 7.6** for an example.

Figure 7.6—Example CCM Control Specification

CCM Control Title	CCM Control Specification	
Business Continuity Management & Operational Resilience Policy	● Establish business continuity and operational resilience policies to ensure appropriate planning, delivery, and support of the organization capabilities in the event of a business disruption. The policies shall align with the organization overall risk management, including its risk appetite.	

7.5.2 Control Audit Frequency

The control audit frequency is used to indicate how frequently each control in the CCM should be audited. For example, some controls may need to be audited annually while other controls may need to be audited quarterly. **Figure 7.7** provides examples of the types of controls that need to be audited for each type of frequency.

Figure 7.7—Example Controls to Be Audited

CCM Control Title	Control Audit Frequency	
Business Continuity Management & Operational Resilience Policy	Annual	
Business Continuity Management & Operational Resilience—Planning	Semiannual	
Business Continuity Management & Operational Resilience—Documentation	Quarterly	

7.5.3 Request List

This section includes a generic list of items that may be requested to support the CCM audit. For each request, there is a section to record who made the request, the date it was requested, who received the request, status of the request and additional notes. **Figure 7.8** includes a list of suggested items to request for the CCM audit; however, this list is not comprehensive—it should serve only as an example. The auditor may need to request additional items during the audit.

Figure 7.8—Request Description

Strategy and Governance
● Provide enterprise cloud strategy and policy. ● Provide enterprise cloud security strategy. ● Provide inventory of third-party attestation reports for the cloud platforms in use. ● Provide the attestation reports themselves and evidence of them being reviewed. ● Provide policies and procedures established for third-party risk assessments, including questionnaires required to assess risk associated with use of third-party services. Provide an example of a completed third-party risk assessment.
Risk Management
● Provide contracts or SLAs of CSPs. ● Provide cloud design documentation, configuration management and provisioning security architecture. ● Provide procedures in place for monitoring control and security processes to provide a secure operating environment. ● Provide CSP's available metrics and controls to assist the organization in implementing the information risk management requirements. ● Provide a list of identified cloud risks, including information on the risk rating and mitigation status.
Compliance
● Provide documentation detailing the legal and regulatory requirements the organization must comply with (e.g., GDPR, PCAOB AS5, PCI DSS, HIPAA). ● Provide documentation related to the implementation of a commonly used risk and controls framework (e.g., ISO, COBIT, NIST) to address regulatory requirements. ● Provide any gap analysis performed against the applicable regulations to determine if there are any regulatory requirements that cannot be satisfied by the cloud computing model.
Third-Party Management
● Provide evidence that a third-party CSP has performed an internal assessment of conformance to its own policies, procedures and availability of control metrics. Provide a copy of the internal assessment of conformance results, if available. Provide a copy of the relevant contract clause that covers this topic. ● Provide a list of all third-party CSPs to the organization for those parts of the organization that are in scope of the audit.
Cloud Data Protection and Lifecycle Management
● Provide data classification and handling policy, including evidence of the last review or update timeline. ● Provide backup and recovery policy, including evidence of the last review or update timeline. Provide latest test results. ● Provide data testing policy, including evidence of the last review or update timeline.
Identity and Access Management
● Provide access provisioning/deprovisioning policies, including procedures implemented for cloud systems and evidence of the last review or update of the policy.
Application Security and Development Operations
● Provide the application design documentation that also governs the subject matter expert involvement in the system design. ● Provide the configuration management and provisioning security architecture.
Incident Response, Notification and Remediation
● Provide the service provider's incident management and response policy and procedures. ● Provide the organization's incident monitoring procedure, which should include both internal and service provider incidents.

Figure 7.8—Request Description *(cont.)*

Threat and Vulnerability Management
• Provide vulnerability management policies and procedures, including evidence of the last review or update timeline.
Key Management
• Provide a listing of the access controls, transmission controls and backup set up to ensure that key stores are in the possession of key managers. • Provide a user access listing of all key store manager's users, which should include their access to all menu options and functions.
Resilience (Backup and Recovery)
• Provide business continuity (BC) and disaster recovery (DR) policy, including evidence of the last review or update timeline. If BC and DR processes are outsourced to a third-party provider, provide the latest third-party assurance reports. Provide evidence that the cloud customer complies with the CSP's business continuity and disaster recovery policies.

7.5.4 Test Result (Pass/Fail)

The next step is intended for auditors to document how they will test whether the organization meets the CCM control. For each test, the auditor marks down whether the organization passed or failed the test. **Figure 7.9** provides one example of a test plan for the business continuity management and operational resilience policy control.

Figure 7.9—Example Test Plan

CCM Control Title	Test Plan
Business Continuity Management & Operational Resilience Policy	• Validate whether the policy has defined roles and responsibilities for business continuity management and operational resilience. • Validate whether roles and responsibilities defined have been properly communicated with the workforce.

7.5.5 Reference to Supporting Documentation

In this section, the auditor includes references to supporting documentation on why and how the organization passed (or failed) the test for the specified CCM control.

7.5.6 Control Audit Conclusion

In this section of the workbook, the auditor adds the final considerations and conclusion, including notes and identifying mitigations, and records the level of effectiveness for control implementation.

7.5.7 Exception Report

Figure 7.10 offers an example template to fill out with any exceptions the auditor identifies during a CCM audit. It includes example observations that are commonly identified while auditing cloud technology management. Note that all exceptions should be carefully assessed and documented. For each exception, the auditor needs to document the following:

- Control test the organization failed
- Root cause of test failure
- Description of exception requested
- Potential risk carried by exception (H/M/L)

- Exception status
- Expiration date

Figure 7.10—Example Exceptions Template

Exception Request				Exception Management	
Control Test Failed	Root Cause of Test Failure	Description of Exception Requested	Potential Risk Carried by Exception (H/M/L)	Exception Status	Expiration Date

7.6 Apply the CCM Auditing Guide (Example—BCR Domain)

Figure 7.11 provides an example of how an auditor applies the CCM Auditing Guide if auditing an organization for the Business Continuity Management & Operational Resilience (BCR) domain. For each control in this section, the auditor starts by writing out the CCM control specification and audit frequency.

Figure 7.11—Applying the CCM Auditing Guide Example

CCM Control Title	CCM Control Specification	Control Audit Frequency	
Business Continuity Management & Operational Resilience Policy	● Establish business continuity and operational resilience policies to ensure appropriate planning, delivery and support of the organization’s capabilities in the event of a business disruption. The policies shall align with the organization’s overall risk management, including its risk appetite.	Annual	

After each control is specified, the auditor needs to document the test conducted on how the organization meets the CCM control. Each control can have one or many test plans, depending on the control being audited. **Figure 7.12** offers examples of test plans for the policy control in the BCR domain.

Figure 7.12—Test Plan Documentation Example

CCM Control Title	Test Plan
Business Continuity Management & Operational Resilience Policy	● Validate whether the policy has defined roles and responsibilities for business continuity management and operational resilience. Validate if roles and responsibilities defined have been properly communicated with the workforce. ● Validate whether the policy is aligned with the organization’s business objectives. ● Validate whether the policy provides a framework for setting business continuity objectives. ● Validate whether the policy includes a commitment to satisfy applicable requirements and continual improvement. ● Validate whether functional and executive leadership reviews the policy on a periodic basis or after significant changes in the organization. ● Validate whether this control is addressed in the shared responsibility model and if the implementation of this model was done correctly.

For each test plan for the control, the auditor then documents whether the organization passed or failed the test, reference any documentation supporting that conclusion, indicate who performed the audit, and list any exceptions (**figure 7.13**).

Figure 7.13—Documenting Test Results

Test Plan	Test Result (Pass/Fail)	Reference to Supporting Documentation	Additional Notes	Audit Performed By	Control Audit Conclusion (Effective/ Exception Noted)
Validate whether the policy has defined roles and responsibilities for business continuity management and operational resilience.					
Validate whether roles and responsibilities defined have been properly communicated with the workforce.					

Figure 7.14 provides an overview of the control audit frequency and test plan for all the controls included in the BCR domain of CCM V4. The auditing frequency in the table is a guide rather than a prescriptive instruction. The auditor should determine the audit frequency of each control based on the nature of the control, the auditee's need for assurance and the requirements of the STAR Program.

Figure 7.14—Control Audit Frequency and Test Plan Overview

CCM Control Title	CCM Control Specification	Control Audit Frequency	Test Plan
Business Continuity Management & Operational Resilience Policy	Establish business continuity and operational resilience policies to ensure appropriate planning, delivery, and support of the organization capabilities in the event of a business disruption. The policies shall align with the organization's overall risk management, including risk appetite.	Annually	Validate whether the policy has defined roles and responsibilities for business continuity management and operational resilience. Validate if roles and responsibilities defined have been properly communicated with the workforce.
			Validate whether the policy is aligned with the organization business objectives.
			Validate whether the policy provides a framework for setting business continuity objectives.
			Validate whether the policy includes a commitment to satisfy applicable requirements and continual improvement.
			Validate whether functional and executive leadership reviews the policy on a periodic basis or after significant changes in the organization.
Business Continuity Management & Operational Resilience—Impact Analysis and Risk Assessment	Determine the impact of business disruptions and risks to establish criteria for developing business continuity and operational resilience strategies and capabilities.	Annually	Validate whether impact analysis includes the following: ● Immediate and ongoing impacts resulting from disruptions ● Maximum tolerable period for disruption ● Estimated internal and external resources required for recovery and resumption

Figure 7.14—Control Audit Frequency and Test Plan Overview *(cont.)*

CCM Control Title	CCM Control Specification	Control Audit Frequency	Test Plan
Business Continuity Management & Operational Resilience—Impact Analysis and Risk Assessment *(cont.)*			Validate whether risk assessment includes the following: ● Identification of critical products and services and their inherent risk ● Likelihood of each risk identified ● Organization risk appetite ● Identification of risk dependencies ● Identification of appropriate and relevant countermeasures to prevent, detect and react to the identified risk
Business Continuity Management & Operational Resilience—Strategy	Establish strategies to reduce impact of, withstand, and recover from business disruptions within acceptable risk appetite. The strategies shall be reflected in business continuity plans.	Annually	Validate whether the strategy is developed by both CSPs and cloud customers within the acceptable limits of their risk appetites.
			Validate whether the strategy covers all aspects of business continuity and resilience planning, taking inputs from the assessed impact and risks, to consider activities for before, during and after a disruption
			Validate whether the strategy covers unavailability of all components required to operate in business-as-usual or in disrupted mode, in part or in total, when impacted by a disruption.
			Validate whether the strategy covers all activities required to continue and recover prioritized activities within identified time frames and agreed capacity, aligned with the organization risk appetite, including the invocation of continuity plans and crisis management capabilities.
			Validate whether the strategy covers all activities within the defined scope to protect the prioritized activities, reduce the likelihood of disruption, shorten the period and limit the impact of disruption on cloud capabilities, and provide for the availability of adequate resources.
			Validate whether the strategy includes detailed solutions and measures for each strategy adopted.
Business Continuity Management & Operational Resilience—Business Continuity Planning	Translate business continuity strategies into consistent, comparable and comprehensive business continuity plans.	Semi-annually	Validate whether the plan is developed consistently in addressing priorities for operational resilience, testing, maintenance and information security requirements.
			Validate whether the plan includes a defined purpose and scope, aligned with relevant dependencies.
			Validate whether the plan is accessible to and understood by those who will use it.

Figure 7.14—Control Audit Frequency and Test Plan Overview *(cont.)*

CCM Control Title	CCM Control Specification	Control Audit Frequency	Test Plan
Business Continuity Management & Operational Resilience—Business Continuity Planning *(cont.)*			Validate whether the plan has assigned roles and responsibilities for review, update, and approval.
			Validate whether the plan has been reviewed at planned intervals or upon significant organizational or environmental changes for effectiveness.
			Validate whether the plan includes a defined communication and escalation plan.
			Validate whether the plan includes a defined method for business continuity invocation.
Business Continuity Management & Operational Resilience—Documentation	Develop, identify and acquire documentation relevant to supporting the business continuity and operational resilience programs. Make the documentation available to authorized stakeholders and review it on a periodic basis.	Quarterly	Validate if the documentation includes a business continuity playbook or user guide.
			Validate if the documentation includes database backup and replication guidelines.
			Validate if the documentation includes architecture diagrams related to business continuity.
			Validate if the documentation includes configuring, installing and deploying changes, and operating the information system.
Business Continuity Management & Operational Resilience—Business Continuity Exercising (including tests)	Exercise and test business continuity and operational resilience plans at planned intervals or upon significant changes.	Quarterly	Validate if the exercises and tests follow processes established in the business continuity plan.
			Validate if the exercises and tests are aligned with business continuity policies.
			Validate if the exercises and tests are performed on critical systems and equipment.
			Validate if the exercises and tests include lessons learned from previous events and exercises.
Business Continuity Management & Operational Resilience—Crisis Management	Establish a crisis management program to develop and apply the adaptive capability of the organization to deal with a crisis.	Semi-annually	Validate if the crisis management program is designed to navigate the organization and its operational teams to predefined thresholds to ensure business recovery.
			Validate if the crisis management program can be invoked and led only by a person of sufficient seniority and experience when the impact of a disruption breaches predefined thresholds.
			Validate if the crisis management program is representative of all key functions necessary to manage a crisis, including authorities and subject matter experts (internal or external).
			Validate if the crisis management program is empowered to ensure the adaptive ability to withstand change as it happens.
			Validate if the crisis management program ensures that all contingency measures, including emergency funds, are made available to the crisis management team when necessary.

Figure 7.14—Control Audit Frequency and Test Plan Overview *(cont.)*

CCM Control Title	CCM Control Specification	Control Audit Frequency	Test Plan
Business Continuity Management & Operational Resilience—Crisis Management *(cont.)*			Validate if the crisis management program ensures that the activities of all teams that are engaged to perform tasks can communicate status at predefined intervals.
			Validate if the crisis management program ensures that sufficient resources are planned to ensure that adaptive spaces, staff, technology and supplies can be acquired when called for.
			Validate if the crisis management program ensures that all criteria to protect the prioritized activities, reduce the likelihood of disruption, shorten the period of disruption, limit the impact of disruption on cloud capabilities, and provide adequate resources are made available to the decision makers.
			Validate if the crisis management program is reflected in policies and business continuity plans and be reviewed and updated when any components change.
Business Continuity Management & Operational Resilience—Recovery	Establish a program to recover any affected business to acceptable levels and eventually to business-as-usual.	Semi-annually	Validate if the recovery program is designed to provide criteria that demonstrate that the impact of the disruption is waning.
			Validate if the recovery program ensures that the operating conditions of the facilities have been adequately reviewed to confirm that they are safe for the people and the technology to operate.
			Validate if the communications to recover businesses are made to the relevant stakeholders.
			Validate if the recovery program has documented processes to restore and return business activities from the temporary measures adopted during and after a disruption
Business Continuity Management & Operational Resilience—Communication	Establish communication with stakeholders and participants during business continuity and resilience procedures.	Semi-annually	Validate that the importance of maintaining effective business continuity and the consequences of failing to do so are communicated to all participants (internal or external).
			Validate that the business continuity and resilience policy, objectives and plans are communicated to all participants.
			Validate that the roles, responsibilities, authorities and expected competencies are communicated to all participants (internal or external).
			Validate that the criteria, thresholds and indicators to demonstrate business continuity are communicated.
			Validate that standard templates for most common communications during a disruption are established for the activation, operation, coordination, and communication of a business continuity response.

Figure 7.14—Control Audit Frequency and Test Plan Overview *(cont.)*

CCM Control Title	CCM Control Specification	Control Audit Frequency	Test Plan
Business Continuity Management & Operational Resilience—Communication *(cont.)*			Validate that a list of the people, technology and processes required for business continuity communications is established.
			Validate that a response structure that will enable timely warning and communication to relevant stakeholders is established.

Page intentionally left blank

7.7 Chapter 7 Knowledge Check

REVIEW QUESTIONS

1. Which of the following considerations does an auditor have to make when performing a CCM-based cloud audit? (Select all that apply.)

   A. In case the organization is in an initial phase of cloud migration, existing baseline security capabilities, tool sets and metrics are irrelevant and potentially misleading.
   B. In case the organization has cloud service implemented, but with immature *ad hoc* processes in place, the auditor should first assess existing cloud services for any significant gaps, and then look at audit procedures which assess security- and governance-related controls.
   C. In case the organization has a formal process in place, CCM auditing is not necessary since the cloud auditing approach can be based on other security requirements from ISO 27001 and SOC 2 TSC.
   D. In case the organization has a formal process in place, the auditor should evaluate the organization against the framework used to assess the third-party services and manage the internal security control framework.

2. There are three recommended steps that an auditor can take to assess the risk level of an organization. Which of the following items belong to those steps? (Select all that apply.)

   A. Evaluate and document key risk conditions
   B. Identify key risk categories
   C. Determine the proper mitigation for each key risk
   D. Establish the organization cloud risk profile

3. Auditors are responsible for scoping the audit and should validate all process and control information provided by the auditee. Which of the following should be documented by the auditor? (Select all that apply.)

   A. Control audit frequency, test plan, test result
   B. Control audit conclusion
   C. Reference to supporting documentation
   D. Who performed audits in the past

4. Consider the following control description, from the Business Continuity Management and Operational Resiliency (BCR) domain:

   Establish business continuity and operational resilience policies to ensure appropriate planning, delivery and support of the organization capabilities in the event of a business disruption. The policies must align with the organization overall risk management, including risk appetite.

   Select the items that should be part of a test plan for this control:

   A. Validate whether the functional and executive leadership reviews the policy on a periodic basis or after significant changes in the organization.
   B. Validate whether the policy provides includes a commitment to satisfy applicable requirements and continual improvement.
   C. Validate the use of automated source code analysis tools to detect security defects in code prior to production which could impact availability.
   D. Validate that the policy establishes the existence of an inventory, documenting data flows for data that is resident (permanent or temporary) within the services' applications and infrastructure network and systems.

Answers on page 344

Chapter 7 ANSWER KEY

Correct answers appear in **bold** font.

1. A. Existing baselines, tools and metrics are always a good starting point for aligning to cloud adoption security risk.

   B. This is the correct approach if the auditee has immature and ad hoc processes in place.

   C. Even if the auditee has a framework already in place, such a framework it has to be anyway cloud relevant.

   D. This is the approach normally followed by an auditor. In addition, depending on the request, the auditor may want to align the CCM controls with the existing controls the organization has defined.

2. **A. This is correct as this is one of the three recommended steps that an auditor can use to assess the risk level of an organization.**

   B. This is correct as this is one of the three recommended steps that an auditor can use to assess the risk level of an organization.

   C. This is incorrect as this is not part of risk assessment but of broader risk management.

   D. This is correct as this is one of the three recommended steps that an auditor can use to assess the risk level of an organization.

3. **A. The information to be documented includes: control audit frequency, test plan, test result, reference to supporting documentation, who performed the audit.**

   B. The information to be documented includes: control audit frequency, test plan, test result, reference to supporting documentation, who performed the audit, control audit conclusion.

   C. The information to be documented includes: control audit frequency, test plan, test result, reference to supporting documentation, who performed the audit, control audit conclusion.

   D. While it is important to leverage the results of previous audits, the information regarding who has performed them is not necessarily relevant when tailoring a new CCM audit.

4. **A. Correct, this test aligns well with the control.**

   B. Correct, this test aligns well with the control.

   C. Incorrect, this is unrelated to the control and the BCR domain in general.

   D. Incorrect, this is unrelated to the control and the BCR domain in general.

Chapter 8:

Continuous Assurance and Compliance

Continuous Assurance and Compliance

Continuous Assurance and Compliance

8.1 Learning Objectives

After completing this chapter, learners will be able to:

1. Explain continuous assurance and compliance.
2. Define DevOps and DevSecOps.
3. Apply DevOps and DevSecOps to security.
4. Outline auditing deployment/CI/CD pipelines.
5. Describe DevSecOps automation and maturity.

8.2 Overview

Automation is key to speed, agility and scalability of the security team. Engineering teams often consider DevOps and automation as synonymous; however, the cultural and organizational changes that DevOps achieves are equally important, but automation is sometimes the more quantifiable benefit.

This chapter provides a basic introduction to DevOps and guidance into specifics of cloud-native development processes and their impact on cloud assessment and auditing. This chapter introduces the concept of automation in the assessment and auditing of cloud processes, and enables the learner to build a mature cloud assessment and auditing process based on automation.

DevOps is highly attuned to the cloud—using DevOps with cloud is a major change in the way organizations have traditionally handled security. The abstraction of cloud, automation and orchestration are very different from how organizations are used to operating. There are two ways to operate, and usually the result is a choice between these two alternatives: native, i.e., organizations that fully adopt the architectures and processes for delivery of services to customers, or tourists, who move existing operational processes and models to the cloud—but when they try to move parts into DevOps, everything breaks.

There is also a third option that combines the transition to a DevOps approach with moving to a public cloud. In a multiyear transition period, teams are individually trained in working with public clouds and adopting DevOps patterns and techniques during a migration period of six to eight months per team. When the organization has a few hundred DevOps teams, the entire transformation takes a few years.

8.3 Concepts of DevOps, DevSecOps

There is a range of definitions for DevOps. Sometimes referred to as a philosophy or a culture, it is a way of breaking down barriers between development and operational teams with the aim of shortening development life cycles and providing continuous delivery with high quality.

DevOps is based on an operational framework that promotes software consistency and standardization through automation. It helps address many difficult development issues related to integration, testing, patching and deployment—by breaking down barriers between different development teams, and by prioritizing tasks that make software development simpler, faster and easier.

DevSecOps requires the integration of security teams and security tools directly into the software development lifecycle, leveraging the automation and efficiencies of DevOps to ensure application security testing occurs in every build cycle. This promotes security and consistency and helps to ensure that security is prioritized no lower than other quality metrics or features. Automated security testing, just like automated application build and deployment, must be assembled with the rest of the infrastructure.

Many of these concepts are either new to security or push security to integrate with development and operations in new ways, using new skills and technologies. In some organizations, they meet with resistance from DevOps teams that have often had a contentious relationship with security in traditionally siloed operations. Stated another way, practitioners of DevOps who have fully embraced the movement may say there is no reason to insert Sec into DevOps, because security is just another ingredient. The DevOps ideal, however, is to break down silos between individual teams (e.g., architecture, development, IT, security and QA) to better promote teamwork and efficiency, and to better incentivize each team toward the same goals. If security is just another set of skills blended into the overall effort of building and delivering software, there is no reason to call it out any more than quality assurance. In practice, that time has not yet come. Establishing security champions in the DevOps teams makes it easier to embed the security function and educate teams on its importance.

However, using the term DevSecOps is a powerful tool for inclusion that serves as an invitation for security professionals to participate as equals. As the movement has evolved, DevSecOps now refers also to leveraging DevOps techniques into traditional security practices. DevSecOps recognizes the differences between the security efforts for coding and the security efforts required for infrastructure and supporting components. The security checks to validate that application code is secure (DevSec) differ from the tooling and processes to verify the supporting infrastructure (SecOps) is secure. These are two different disciplines, with different tooling and approaches, and it is a mistake to discuss one as a subset of the other.

This is not about security getting control over development. DevOps is about working alongside partners, not controlling them. This section provides some proven examples of how to collaborate with development teams in a positive-sum way.

There is no one strict way to do DevOps, but it always includes three overlapping elements:

- Continuous integration
- Continuous delivery
- Continuous deployment

Continuous integration manages code building and testing, while continuous delivery refers to the ongoing creation of new features or artifacts. Continuous deployment refers to the assembly of the entire application stack into an executable environment. Opportunities for improvement are in each element of the DevOps process, resulting in an easier, faster and more reliable development life cycle. Moving to DevOps can be a time-intensive change. Building out the basic scripts and templates can take months; maturing them into a reliable software delivery infrastructure can take years.

8.3.1 Continuous Integration

Continuous integration (CI) allows developers to regularly check in small iterative code advances in new builds or updates to a shared source code repository, often daily. The key is to make smaller, simpler additions, which makes finding code defects easy and quick.

DevOps implements Agile concepts in processes that drive code. In DevOps, CI implies that code is not only built and integrated with supporting libraries, but also automatically tested. DevOps Continuous Integration works with many different development and integration models, including integrating changes into branches and the main body of the code, which reduces the complexity and integration problems that can plague development teams.

DevOps requires considerable supporting infrastructure. Builds must be fully scripted, and the build process occurs as code changes are made. After every successful build, testing receives the bundled application stack. Before unit, functional, regression and security testing, test code is built and checked into repositories as part of the automated process.

When a new application bundle is available, tests automatically begin; therefore, new tests are applied with each new build. Test systems must be automatically provisioned, configured and populated with data, prior to tests being launched. Automation scripts must monitor the whole process and communicate results to development and operations teams as events occur. Operations, testing and development teams work together closely to create the scripts and tools.

Figure 8.1 shows the pipeline for an automated build, which includes security test points for containers. The pipeline evolved, through continuous improvement, to reach this level of orchestration.

Figure 8.1—An Automated Build Pipeline

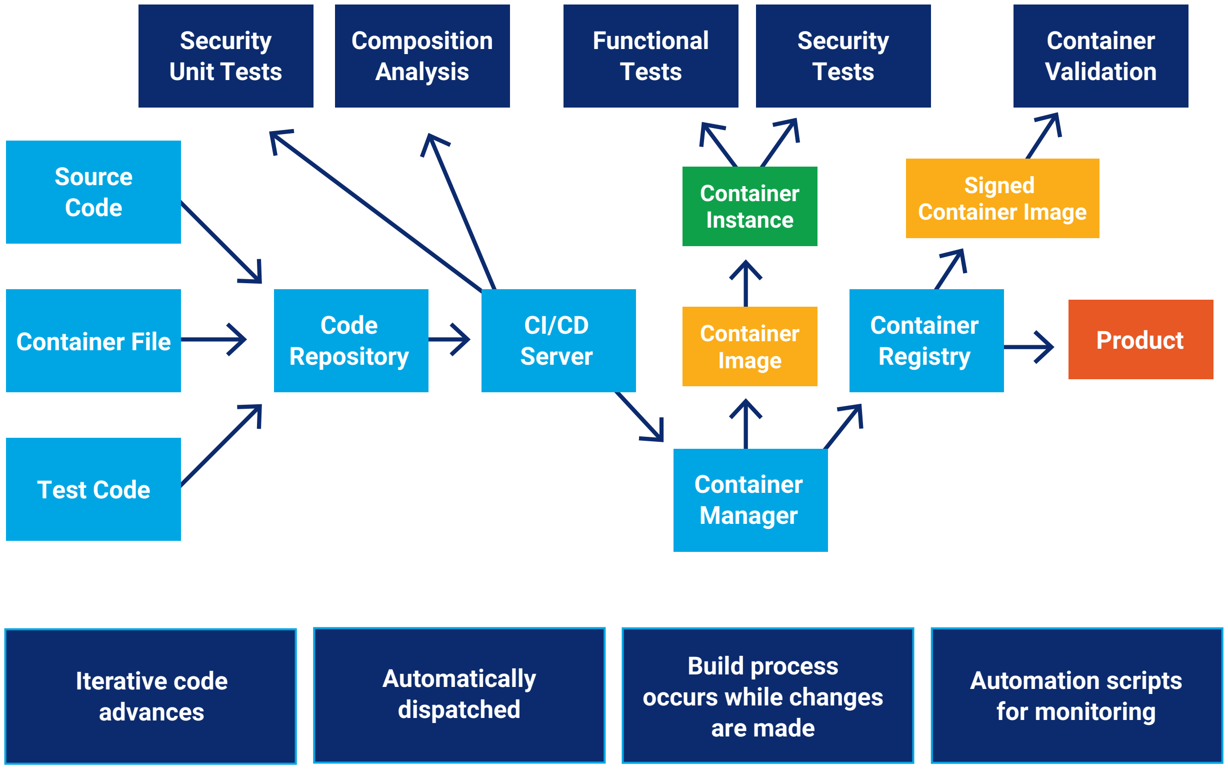

Source: Securosis

8.3.2 Continuous Deployment

The packaging, testing and monitoring tasks of continuous deployment (CD) are very similar to CI but with the purpose of releasing software to end users rather than building it. The results of a successful build cycle feed the CD process. CD automates the release management, provisioning and final configuration for the application stack, and then launches the new application code.

Organizations can employ the CD concept in one of two ways. An organization can launch a new version of its application into an existing production environment where the application layer is automated, but the server, data and network layers are not. On-premises applications and private cloud deployments often use this model; some public cloud deployments continue to use this model. The other model uses infrastructure configurations (see **figure 8.1**) that can be a simple Dockerfile or a complex, large CloudFormation or Terraform® package to build out an entire cloud infrastructure.

Many security teams are apprehensive about or mistrust CD, believing it allows code to go live without checks or oversight; but, with CD, the code does not go live until it passes all security tests, and some CI pipeline tests contain manual inspection points. Experienced secure code specialists say that CD results in more reliable and more secure releases, due to the way it eliminates many issues that are inherent to code deployment, including error-prone manual changes and discrepancies between production and development in supporting-library revisions. When sufficient testing is in place, there should be no reason to mistrust CD.

8.3.3 Managed/Blue-Green Deployments

Most organizations follow the managed release approach, which uses manual execution and timing, and automated actions, including automated infrastructure deployments. These organizations have a slower release cycle, often going live after one to three sprints. Managed releases may cycle the entire infrastructure stack with the application. This deployment type relies on automating the environment the software runs in. This can be done by standing up a Kubernetes® cluster, leveraging Openshift® to run Terraform templates into Google Cloud Platform™, or launching an entire cloud environment via Infrastructure as Code templates. The infrastructure and the application are all code, so they are launched in tandem. This approach is becoming common in public cloud deployments.

The managed release method offers significant advantages and provides the foundation for blue-green or red-black deployment—old and new code will run side by side, as close mirror images, with each on its own execution environment. The old (blue) code continues to serve user requests; the new (green) code is exercised only by select users or test harnesses. A rollout is a simple redirection at the load balancer level. Internal users and live customers can be redirected gradually to the green servers, and function like testers for the new system.

If the new code passes all required tests, the load balancer sends all traffic to the green servers, the blue servers are retired, and then green becomes the new blue. If errors are discovered, the load balancers are pointed back to the old blue code until a new patched version is available. This method is performing prerelease testing in production and near-instantaneous rollback if defects or security problems are discovered.

8.3.4 Integrating Security Into Deployment Pipelines

Wide use of automated application security techniques within the code development process is relatively new. In the past, application security was often added with network or application firewalls, and not a part of the code itself. Security product vendors discovered that it was very difficult to understand application requests from outside the application to detect and block attacks. Fixing vulnerable code and closing off attack vectors when possible is much more effective. Although add-on tools are improving (some work inside the application context), addressing the issues in the code is more effective.

Shift left is a central concept of building security into development that integrates security testing earlier, within the software development life cycle (SDLC). The phases of the SDLC are typically listed left to right—design, development, testing, preproduction and production. With shift left, more resources shift away from production on the extreme right, and more resources go into the design, testing and development phases, on the left. This concept originated in lean manufacturing, in Kaizen and Deming's principles. Although effective, these ideas typically were limited to the manufacture of physical goods. DevOps promoted their use in software development and demonstrated that security can be improved at a lower cost, by shifting security defect detection to earlier in the process.

8.3.5 Models/Methods of CI/CD: Commits, PRs, Patterns Like Git Flow

With a CI pipeline (**figure 8.2**), the left side can have different types of text files, including source code, infrastructure templates (e.g., CloudFormation), and then server or container definitions that will configure them and make sure the right operating system is in use. Any type of application stack can be defined in those files.

Figure 8.2—Security Checking Pull Requests

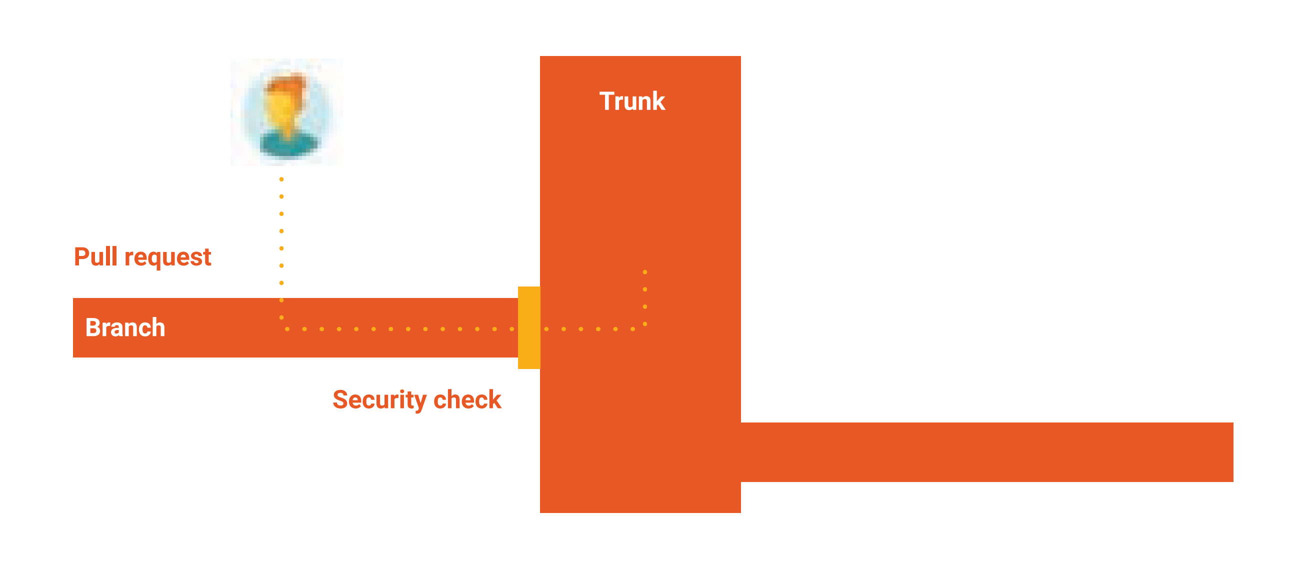

The text files are kept in version control repositories that are typically different versions of Git. When someone modifies them, they synchronize locally and are fetched into a new commit. These commits can be handled in several ways. In larger, more complicated projects, a developer might work on one part, clone a copy, perform the work and commit those changes in the branch. Those changes then need to be merged into the trunk. Changes to the branch or trunk submitted for approval are known as pull requests (PRs).

A PR is used when a developer suggests a change to a repository that someone else owns. The owner can review the change and decide whether to merge it. The commit is a change, and the PR is a request to change that file. Not every Git system accepts PRs, and some enterprises don’t use PRs but use straight commits instead.

It is important for security professionals to have a basic understanding of the development process. They need to understand the terms and the DevOps team’s chosen methodology, because that determines where they do their security checks. Some teams do not do security testing on every commit—they do it on the PR for efficiency reasons, because some tests can take longer to run. There is no one right way to do it. Organizations may make the same tool available to run optionally on commits, but testing then becomes mandatory before anything is added to the main feature branch. It is important for the auditor to understand the flow in the process, to determine that those tests are run before the code is moved into production, and to know where the security gates are.

Git flow, for example, is one pattern for organizational management of commits and PRs. Other enterprises use entirely different models based on the same core technology and concepts. It is important to understand what pattern is in use, because it defines optimal positioning of security and audit gates. What is the pathway or pattern the organization uses to produce artifacts—the actual software? What is its flow? Workflows are like an assembly line floor map that defines where teams place security tools and tests for optimal effectiveness.

8.3.6 Leveraging Automation for Cloud Security Operations

Like Agile, DevOps uses automation and more frequent software releases containing fewer code changes in each, to bring speed, consistency and efficiency to development, operations and testing. Automation and fewer changes per release allows for better focus and improved clarity of purpose with each release, resulting in fewer mistakes. Rolling back to a previous release to fix mistakes is easier with DevOps. Because automation software repeats the same processes every time, consistency is the most conspicuous benefit.

The application build servers are where automation is first applied and its benefits of are most pronounced. Build servers (e.g., Bamboo, Jenkins®, CircleCI℠), also called CI servers, automatically construct an application—and possibly the entire application stack—as code is changed. Once the application is built, these platforms may also launch quality assurance (QA) and security tests, kicking failed builds back to the development team.

Automation benefits other facets of software production, including reporting, metrics, quality assurance and release management. It is also possible to automate many aspects of security testing to improve the overall security of code and other artifacts. At the outset, this may not seem like much, i.e., calling security testing tools instead of manually running the tests. That perspective misses the fundamental benefits of automated security testing. Automation ensures that each update to software includes security tests, promoting consistency. Automation also promotes efficiency by avoiding mistakes and omissions common with repetitive manual tasks. But most importantly, as security teams are typically outnumbered by developers, automation is the key ingredient for scaling security coverage without having to scale security personnel headcount.

Automation is used to improve overall security operations unrelated to development. Guardrails, for example, are one form of automation that can remediate misconfigurations in cloud deployments. Many organizations are rapidly integrating automation into their incident response playbooks. Although not all of these are directly DevSecOps, they typically involve leveraging the same automation tools and technologies used by traditional DevOps, but in the security context.

8.4 Auditing Deployment/CI/CD Pipelines (Part 1)

DevOps, including both pipelines and automation, creates the opportunity to improve compliance and auditability by enforcing consistency and keeping strong artifacts (logs) of activity. Essentially every change is tracked and attributable to the submitter, and automation ensures consistent application of policies. This affects audits, since most of the focus is on ensuring that pipelines and automation are properly configured and secured, and on applying the expected policies and controls. Instead of compliance being caught in an endless assess/remediate cycle, compliance is enforced continuously using policies (code). This is what is meant by terms such as continuous audit, continuous compliance, and compliance as code.

A well-structured CI/CD pipeline offers full auditability and attribution for any change that touches any environment, including production. A point worth noting is that, assuming the organization enables all audit logs, it's possible to trace every single change back to the developer or administrator who approved it, and which security tests it passed or failed. From an audit perspective, this is an extremely positive development. An assessor will want to focus on evaluating the inherent risk security of the pipelines:

- Are they patched and up to date?
- Are they protected on the internal network?
- Are there proper controls for access to the pipeline?
- Does the pipeline provide proper artifacts for audit (and are they enabled by default)?

The auditor should make sure the pipeline is secure, and that the right auditing capabilities are turned on within the pipeline, and then make sure that the security gates (security tests that must be passed to proceed in the process) or testing are properly implemented in the pipelines. Audit logs should include the following:

- Modifications to the CI/CD configuration (e.g., installing a plugin)
- Modification of job settings
- Usage of the pipeline: starting jobs, stopping jobs, manual approvals, etc.
- Changes to RBAC, and assigning permissions and credentials

The keys for security practitioners are twofold: Security policies can also be defined in scripts/code, and that way it's possible to examine code repositories to ensure the templates, scripts and code are secure *before they run*. This is a fundamental change to security audits.

8.4.1 Types of Integrated Security Tests

After ensuring the pipeline itself is secure and effective artifacts are provided, the next step is to evaluate the security testing and gates. Based on the flow pattern and technical architecture of the pipeline, what are the proper ways to perform security tests? Which tests fail builds, and which create issues for later remediation? Who defines and maintains the tests and interprets the results, and are those tests sufficient?

A good way to think of this is in terms of paths and gates. The path is the flow from a developer or administrator making a change to the update being deployed. It starts on the administrator's computer and flows through the various version control repositories, CI/CD servers, and other technologies used to create, test and deploy the software.

Any given pipeline will likely have different paths for the various kinds of artifacts and changes. For example, code built into containers will follow a different path than code pushed into virtual machines, which will be different from pushed infrastructure changes. The goal in an audit is to understand the paths and then look for the gates. Where are security tests performed? What kind of tests? Is the test scheme complete and effective? Which gates stop the process and which gates create issues to fix later? What are the processes for tracking issues over time?

The goal is to ensure that there are security gates along every path and that every change is tested and audited (**figure 8.3**).

Following are the most common kinds of software security tests. There are technical limitations on where these tests can be performed, often based on timing. For example, static application security testing (SAST) may need to run overnight on large code bases, and it usually results in identifying issues to be remediated later. High-level exposures, like a failed credential scan, should always block the process; the tests can run quickly.

Figure 8.3—Using Security Gates

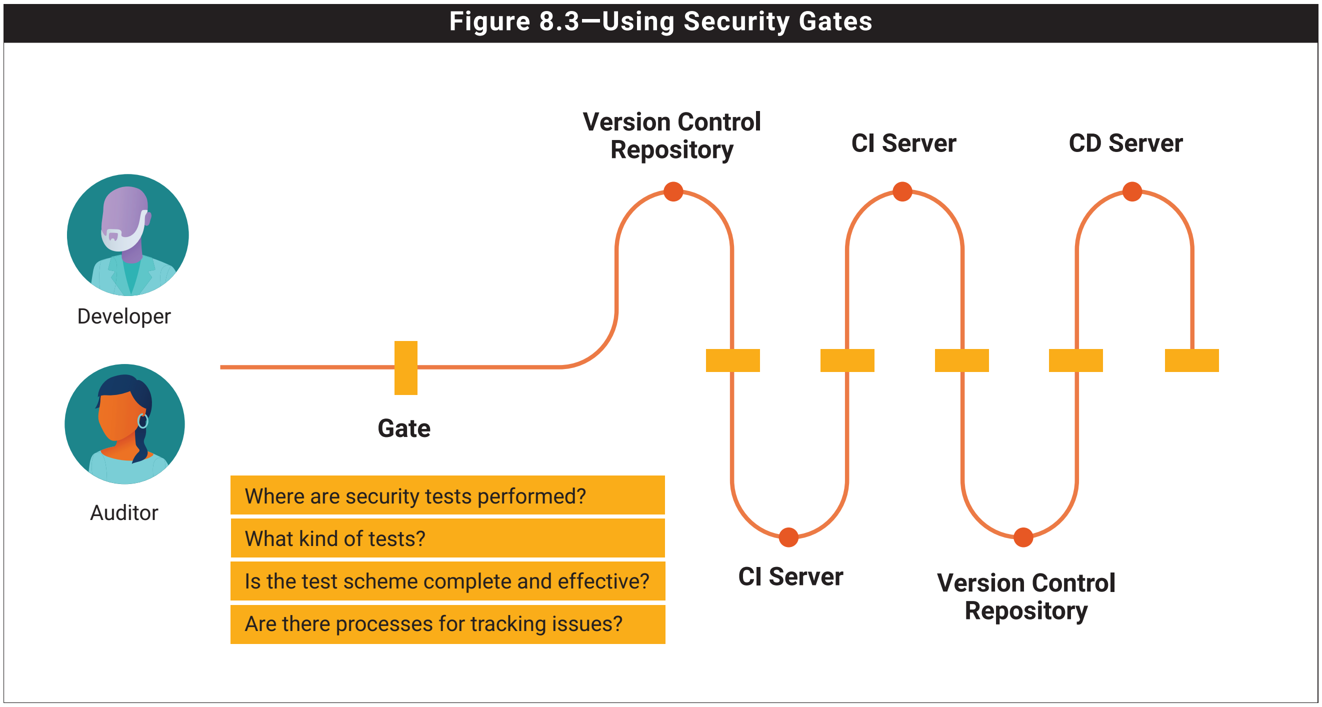

Where are security tests performed?
What kind of tests?
Is the test scheme complete and effective?
Are there processes for tracking issues?

SAST

Static application security testing (SAST) examines all code—or runtime binaries—to support a thorough search for common vulnerabilities. SAST tools are highly effective at finding flaws, even in code that has been manually reviewed. The selection criteria will likely boil down to speed of scan, ease of integrations, readability of results, and lack of false positives. Most of these platforms have become effective at providing analysis that is useful for developers, not just security professionals.

Many of the products are being updated to offer full functionality via APIs or build scripts. If there is a choice, select tools with APIs for integration into the DevOps process, and those that do not require "code complete." There has been a slight reduction in use of these tests—as they often take hours or days to run—in a DevOps environment that can prevent them from running inline as a gate to certification or deployment. Most teams are adjusting to support out-of-band—or what is sometimes called "parallelized"—testing for static analysis. It is highly recommended to keep SAST testing inline if possible and focus on new sections of code to reduce runtime.

Stored Credentials/Secrets Scanning

A frequently occurring problem in security is the storage of credentials such as passwords and access keys in source code, or configuration or other files. This creates the risk that those credentials—which should never be shared—are exposed to people who do not have permission to access them. This could lead to the organization's application or data being vulnerable. Credential scanning tools search source code for credentials to make sure they are not inadvertently checked in with code.

Security Unit, Functional and Regression Testing

DevOps encourages testing in all phases of development and deployment. The security tests typically sit side-by-side with functional and regression tests that quality assurance teams have likely already deployed (**figure 8.4**). Beyond those typical post-build testing points are testing on a developer's desktop prior to check-in, in the code repositories before and after builds, and in predeployment staging areas.

Figure 8.4—Unit, Functional and Regression Testing in DevOps

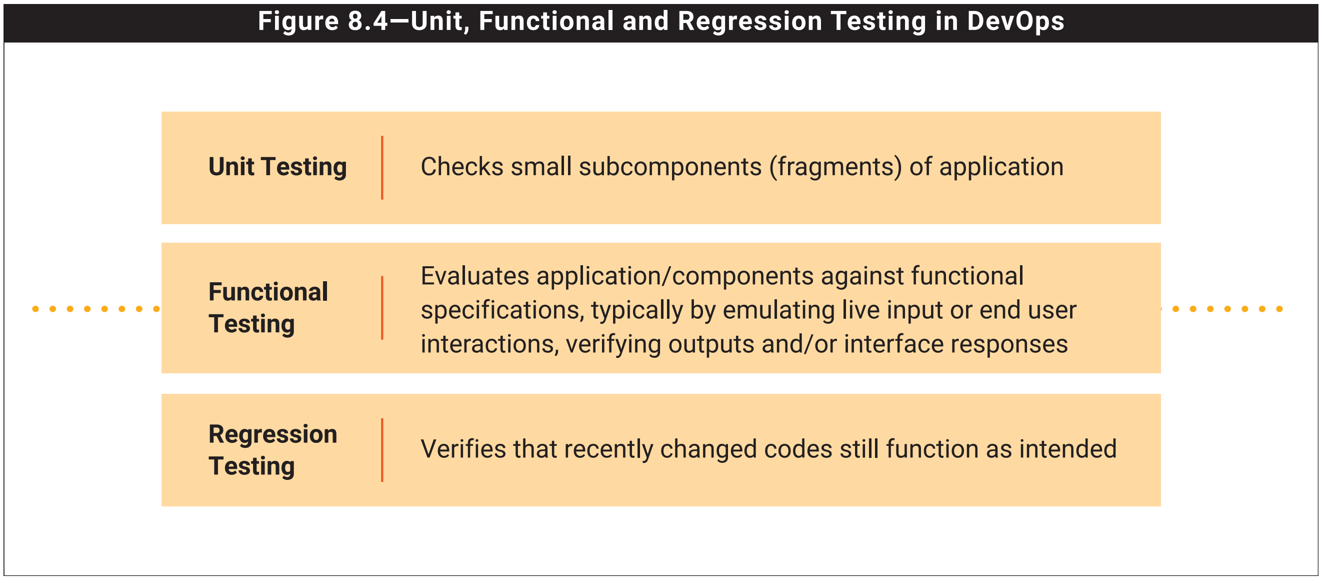

Unit Testing	Checks small subcomponents (fragments) of application
Functional Testing	Evaluates application/components against functional specifications, typically by emulating live input or end user interactions, verifying outputs and/or interface responses
Regression Testing	Verifies that recently changed codes still function as intended

Unit testing is checking of small sub-components or fragments (units) of an application. These tests are written by programmers as they develop new functions and are commonly run by developers prior to code check-in, but possibly within the build pipeline as well. These tests are intended to be long-lived, checked into the source repository along with new code, and run by every subsequent developer who contributes to that code module. For security these may be straightforward—such as SQL injection against a web form—or more sophisticated attacks specific to the function under test, such as business logic attacks—all to ensure that each new bit of code correctly reflects the developer's intent.

Every unit test focuses on specific pieces of code, not on systems or transactions. Unit tests attempt to catch errors very early in the process, per Deming's assertion[259] that the earlier flaws are identified, the less expensive and easier they are to fix. In building out unit tests, it is beneficial to support developer infrastructure to design effective tests, and to encourage the team to take testing seriously enough to build good tests. Having multiple team members contribute to the same code, with each writing unit tests, helps identify weaknesses a single programmer might not consider.

A regression test verifies that recently changed code still functions as intended. In a security context, this is particularly important to ensure that vulnerabilities are not reintroduced. DevOps regression tests are commonly run parallel to functional tests—after the code stack is built out. In some cases this may need to be in a dedicated environment, because security testing can be destructive and cause side effects that are unacceptable in production servers with real customer data.

Virtualization and cloud infrastructure are leveraged to expedite instantiation of new test environments. The tests themselves are typically home-built test cases to exploit previously discovered vulnerabilities, either as unit or systemic tests. These types of tests ensure that credentials like passwords and certificates are not included in files, and that infrastructure does not allow port 22 or 3389 access.

[259] The Deming Institute, "Dr. Deming's 14 Points for Management," https://deming.org/explore/fourteen-points/

Software Composition Analysis (SCA) and Vulnerability Analysis

Software composition analysis (SCA) tools check versions of open-source libraries to assess open-source risk, both for security vulnerabilities and potential licensing issues. Some people equate vulnerability testing with DAST (dynamic application security testing—see paragraph below), but there are other ways to identify vulnerabilities. In fact, there are several kinds of vulnerability scans. Some look at settings like platform configuration, patch levels or application composition to detect known vulnerabilities. Issues like Heartbleed, misconfigured databases, and Struts vulnerabilities may not be part of an organization's application testing at all, but they are critical application stack vulnerabilities. Scans should be broadened to include the application, application stack, and the open-source platforms that support it.

DAST

Rather than scanning code or binaries like SAST, dynamic application security testing (DAST) dynamically crawls through an application's interface, testing how it reacts to various inputs. DAST scanners cannot see what's going on behind the scenes, but they offer valuable insight into how code behaves, and they can flush out errors other tests may not see in dynamic code paths. The good news is they tend to have low rates of false positives. These tests are typically run against fully-built applications and they can be destructive, so the tools often include settings to vary between test and production environments. Like SAST, DAST may require some time to fully scan code, so in-line tests that gate a release are often run against new code only, and full application sweeps are run in parallel.

Assessing Infrastructure as Code

Applications are software. This is fairly well understood, but less appreciated is that in many environments—especially public cloud—servers, networks, messaging, IAM and every other bit of the infrastructure may be defined as configuration scripts, templates or application code. IT teams now define entire data centers with templates comprising a few hundred lines of script.

Infrastructure as Code templates need to be assessed like any other code, but tooling is still in the earliest stages. Standard SAST tools are designed to work with programming languages, not the template formats used by these tools. However, open source and commercial tool options are emerging and being successfully integrated into deployment pipelines.

Since the same templates used to build production environments can often be used to build test environments, an alternative option is to deploy the templates into tests and then rely on DAST, configuration assessment tools, or even manual acceptance testing to validate that security and compliance requirements are in place.

The ideal is to pick up misconfigurations and vulnerabilities before the code deploys, but combining static analysis and dynamic testing of a deployed template is currently the most viable option.

8.5 Auditing Deployment/CI/CD Pipelines (Part 2)

8.5.1 Testing Locations

The key point for auditors is to verify that the security gate is somewhere on the path; it's less important where it is located. This is about making sure the protection is in place, and that nothing goes through the pipelines in each of those tests—it's not necessary to be prescriptive about where the tests occur.

A common problem to all Agile development approaches is what to do about tests that take longer than a development cycle. For example, fuzz testing critical pieces of code takes longer than an average Agile sprint. SAST scans of large bodies of code often take an order of magnitude longer than the build process. DevOps is no different—with CI and CD, code may be delivered to users within hours of its creation, and it may not be possible to perform complete static analysis or dynamic code scanning. To address this issue, DevOps teams run multiple security tests in parallel to avoid delays. They break down large applications into services to speed up scans as well. Validation against known critical issues is handled by unit tests for quick spot checks, with failures kicking code back to the development team. Code scanners are typically run in parallel with unit or other functional tests.

Security professionals should look for ways to speed up security testing. Organizing tests for efficiency versus speed—and completeness versus time to completion—is an ongoing balancing act for every development team. Focusing scans on specific areas of code helps find issues faster.

One approach to avoid latency is to maintain and fully configure test servers—just as with production servers—waiting for the next test cycle. Rewriting and reconfiguring test environments for efficiency and quick deployments also helps with CI.

Version Control Repository

Source code management, configuration management databases, container registries and similar types of tools store code and help with management tasks like versioning, IAM and approval processes. From a security standpoint, several composition analysis vendors have integrated their products to check that open-source versioning is correct and that the platforms in use do not contain known common vulnerabilities and exposures (CVEs). Additional facilities to create digital fingerprints for known good versions and other version control features are becoming more common.

CI/CD Server

Software development leverages many tools to manage code and process. The two most important to security are code repositories and build tools. Repositories like Git essentially manage application code, giving developers a shared location to store code, and track versions and changes. Others, like Docker Registry, are specifically for containers. These tools are essential for developers to manage the code they are building, but they are also important to security as a place where code can be inspected. Build servers like Jenkins and Bamboo automate the building, testing and delivery of code. Rather than at a component or module level, they are typically employed for full application stack testing. Developers and quality assurance teams use the build server to launch functional, regression and unit testing. Security teams should leverage these build servers to integrate security testing (such as SAST, DAST, composition analysis, and security unit tests referred to in the previous section) so it fits within the same build process and uses all the same management and communications tools. It is important for security teams to understand which tools the development team uses and who controls those resources, and to arrange for integration with security testing.

Artifact Storage

Build servers enable the creation, testing and deployment of applications. Often comprising multiple sources of in-house-developed and open-source code, it is common for many artifacts to be used in the construction of an application. Build servers like Jenkins and Bamboo have the hooks needed to massage these artifacts, both before and after a build. This is commonly how testing is integrated into the build pipeline. Leveraging this capability is central to security test integration. Composition analysis, SAST and custom tests can all be integrated at this stage.

Dev/Test/Prod Environments

With production environments more locked down, development and test servers become more attractive targets. Traditionally these environments ran with little or no security, but the need for secure source code management, build servers and deployment pipelines is growing in light of the importance and sensitivity of code.

For “code complete” or systemic testing where all parts of the application and the supporting application stack are assembled, a preproduction staging area should be set up to mimic the production environment and facilitate a full, more accurate battery of tests. This kind of preproduction testing is becoming a notable trend of late. As public cloud resources allow for rapid elasticity and on-demand resource procurement, firms are spinning up test environments, running their QA and security tests, and then shutting them down again to reduce costs. In some cases, this is used to do DAST testing that used to be performed in production. And in most cases, this means the previously described blue-green deployment model is leveraged to run many different types of testing on a new environment parallel to the existing production environment.

8.6 Auditing Deployment/CI/CD Pipelines (Part 3)

8.6.1 Auditing the Security and Compliance of Pipelines

To approach auditing the security and compliance of DevOps pipelines, auditors should start with an architectural assessment because there are many different approaches. One is to map out where the security gates are and how the pipelines touch production and development environments (e.g., where secrets management is located, where credentials are stored—that is all part of architecture).

Because CI/CD pipelines offer an automated pathway into production, stricter access controls should be implemented for these systems, particularly build servers and code repositories. And because scripts run continuously in the background with minimal human oversight, an additional monitoring process or control is needed to catch errors and misuse.

8.6.2 Pipeline and Secrets Management

Many pipeline tools offer good security, with digital fingerprinting, MFA, logging, RBAC and other security features. When deployed in cloud environments, where the management plane allows preventive control of an entire environment, great care should be taken with access controls and segregation of duties. There should be RBAC plugins and some kind of secrets manager. Depending on the pipeline, there should also be plug-ins to audit job and configuration changes.

Secrets managers are dedicated tools for handling secrets such as passwords, access keys, and tokens in a secure manner. Aside from serving as hardened vaults, they typically include capabilities to rotate secrets, generate single-use credentials, and overall reduce or eliminate stored static credentials that pose a huge security risk. They are especially useful for pipelines that require deep access to carry out their tasks but may not have the best inherent security.

8.6.3 Separation of Development and Production Pipelines

Continuous integration servers or pipelines have a lot of credentials because they need to be able to make changes in a production environment. Best practice is to separate the pipeline that does the production pushes from the development pipeline. This reduces the number of people who can touch that production environment through the pipeline, and it isolates developers from directly modifying production.

For practical purposes, it's acceptable to have a shared version control repository and a shared artifact repository (e.g., the compiled source code). However, the automation that moves this code into a production environment should remain separate.

Thus, the two key components to keep separated and isolated are the CI/CD server/tool itself, and any credential vaulting/secrets manager used to hold the production and development credentials. For example, an organization can use a shared version control repository and a shared artifact repository. Then there is a separate CI/CD pipeline for production that pulls from the same code and same artifacts but uses a separate secrets manager that holds the production credentials. After a build is approved in development and passes all its tests, that build is marked for promotion to production. Then—either automatically or manually by a release manager or similar authority—the production build is pushed.

Some organizations use a security ticket system for tracking security testing in development. In production, instead of repeating security tests, the tickets are checked. Then the code and other artifacts are validated as being those that passed the testing, and the build is promoted (this is largely automated).

8.6.4 Impacts of Containers and FaaS/Serverless

Containers and serverless computing are different from traditional environments. The auditor should be familiar with the underlying technologies to carry out proper assessments. Containers always have containment and orchestration levels (e.g., Kubernetes), and the configuration of the container management system should be within scope of assessments.

Containers impact security and compliance in a few ways:

- The container orchestration system adds a new management plane, even when running on a cloud platform. This is subject to extensive security and compliance implications, including fundamental security (e.g., is it patched, up-to-date and configured securely?) and IAM/access (e.g., who is authorized to access it and make changes?). These systems lack the isolation controls of major cloud platforms, and orchestration systems need to be examined closely to ensure security boundaries are enforceable.
- Containers are defined in code (templates) and create artifacts (binaries) that are stored in repositories and registries that must be properly configured and secured.
- Which containers are used where has significant compliance implications. In audits, it's important to ensure that only the authorized containers run on their designated hosts/clusters, and that security context and rules are enforced.
- Containers typically go hand-in-hand with service meshes, which are a further automation and management layer used to connect application components. For compliance and audit, this means configuration, security boundaries and authorizations are within scope.
- Containers include parts of the operating system, software and code that need to be audited and assessed like any other server or virtual machine.
- Serverless (Function as a Service) is a newer technology that further abstracts application code and removes container and server configurations from the picture. FaaS always runs in a container on some server, but it is abstracted and opaque to the user. Key compliance factors include ensuring the following:
  - The code still runs through all the appropriate security testing.
  - The FaaS configuration meets requirements (e.g., use of encrypted environment variables or network controls).
  - The IAM of the serverless functions is properly configured with least privilege applied.

8.7 DevSecOps, Automation and Maturity

In cloud and DevSecOps, more mature environments are more consistently and effectively managed, with fewer security defects exposed in production environments, thus reducing attack surface and risks. It's possible to gauge the maturity of an environment based on how much automation is present. In mature environments, the audit effort can focus more on the pipeline and security gates, while in less mature environments the audit will have to rely more on technical assessments of the deployed assets.

Note: Key to measuring maturity is whether any of the security testing is automated, security testing is performed in parallel, or test results are acted upon. The more security processes an organization can automate and embed, the faster it can move, while simultaneously reducing risk. **The key focus: In the end, is the organization producing better code, and is it able to do it faster and consistently?**

In less mature organizations, testing may find many flaws, but they are not fixed in a timely or consistent manner. There may also be silos between teams and a high degree of reliance on manual assessments and remediations. Or there may be security rules that are inconsistent with cloud and DevOps practices, which impedes the process without providing any additional security benefit.

Note: The key indicators to assess:

- How well do the development and security teams work together?
- Are the security controls and testing aligned with the organization's objectives, compliance requirements and risk tolerance?
- Do the pipelines and automations consistently produce code with better auditability and fewer defects and compliance findings?

8.7.1 Auditing DevSecOps Process Design and Implementation

In a DevSecOps context, audits start by understanding the patterns and processes the organization uses. These will define where security testing and gates should be deployed to maximize their overall effectiveness. This, in turn, closely correlates to maturity. More mature organizations have greater degrees of security automation and consistently produce more secure applications, including the supporting infrastructure, thanks to using Infrastructure as Code.

Note: The start of any audit is to map out the overall operational processes used to produce code, including the processes for security testing and the positioning of security gates.

Mature organizations, in particular, may have multiple approaches with different levels of maturity, but for an auditor the objective is to ensure that no matter the process, the deployed artifacts go through required security testing and audit log creation, and the organization has supporting processes to act on findings. When implemented properly, with gates and the proper test definitions, this ensures continuous assurance and compliance in rapid DevSecOps deployment cadences.

Pipeline Assessments

The key component to continuous assurance and compliance is embedding security testing and audit logs throughout the pipeline. This includes three primary components:

- Security tests/gates where the tests are performed

- Proper security and compliance tests performed at each gate
- Effective audit log generation and collection for test results, and change management for both deployments and the pipeline itself.

Architectural Assessments

Reference security architectures exist for different types of applications and services, including web applications, data processing applications, identity and access management services for applications, stream/event processing and messaging.

The architectures are even more effective in public cloud environments, Kubernetes clusters, and service mesh environments—where it is possible to tightly control via policy how each application operates and communicates. With cloud services, it is recommended to leverage service provider guidelines on deployment security, and while they may not be called "reference security architectures," they are offered. Assessors should educate themselves on the application platforms and ask software designers and architects which methods they employ. It should not be surprising if there is no direct answer for legacy applications, but new applications should include plans for process isolation, segregation and data security, with a full IAM model to promote segregation of duties and data access control.

Operational Standards

Collaboration with development teams is essential to define minimal security testing requirements, and critical and high-priority issues. Assessors will need to negotiate which security flaws will fail a build and define the process in advance. There probably should be agreement on time frames for fixing issues, and some type of virtual patching to address hard-to-fix application security issues. These things should be defined up front and the assessor should make sure development and IT partners agree.

Manual Assessments

Many enterprises find it somewhat scary to fully automate deployment, so they want a human to review changes before new code goes live. There are also very good security reasons for manual review. In an environment as automation-centric as DevOps, it may seem antithetical to use or endorse manual code reviews or security inspection, but manual review is still highly desirable because some types of vulnerabilities are not captured by scanning tools. Manual reviews often catch obvious things that tests miss, and that developers can miss on their first (only) pass. Also, developers' ability to write security unit tests varies. Whether through developer error or reviewer skill, people writing tests miss defects that manual inspections catch. The tool belt should include manual code inspection, at least periodic spot checks of new code, or things like Dockerfile, which are often omitted from scans.

Automated Assessments

Automation is how security analysis can be faster, more frequent, and without direct involvement from the security team. Security tools that perform automated analysis, either out-of-the-box or custom checks, are critical to scale across development teams, and the organization should integrate them into every build pipeline. But this means not only that scans are automated, but also that distribution of results is integrated with other tools and processes. This is how teams scale.

Remember that development is not the only group writing code and building scripts; operations is now up to its elbows as well. This is how DevOps helps bring patching and hardening to a new level. Operations' DevOps role is to provide scripts that build out the infrastructure for development, testing and production servers. The good news is that this is now testing exact copies of production. Templates and configuration management address a problem traditional IT has struggled with for years—*ad hoc* undocumented work that tweaks the environment to get it working. There is a great deal of work to get environments fully automated—on servers, network configuration, applications and so on—but it makes future efforts faster and more consistent. Many teams build new machine images every week, update their scripts to apply patches, and update their configurations and build scripts for different environments. In other organizations, the central cloud infrastructure team makes new images available. Combined with auto patching, this ensures consistency and a secure baseline.

8.7.2 Continuous Assessment Tooling

With the proper assessment process in place, the next step is to leverage the tooling for continuous assurance and compliance. The testing should first be put in place, and then the processes and additional tooling should be built to interpret and manage the results.

These key tools and indicators enable a continuous assurance DevSecOps program:

- Major architectural design changes and approvals can be kept in either the version control repository or the issue tracker (or both). Mature DevOps organizations use issue tracker tooling (various Kanban-board style products) for defining their software development epics and sprints. This creates a record of design decisions to create accountability at the design phase. DevSecOps organizations can include security and compliance sign-offs within this process and the tooling will maintain needed records.
- Deployment pipelines should keep two specific audit log trails to ensure the pipeline itself remains compliant. The first set is for any structural changes in the pipeline, such as software updates, including updates to the underlying host/platform. All administrative logins should be recorded for pipeline host operating systems or other services. The second set is for changes to build jobs (e.g., changing the flow of the build, adding new kinds of tests, or removing tests). Changes to either the pipeline or build structure should trigger manual review.
- There are three major sources to track within the pipeline itself. These results should generate the following as part of the software development process:
  - The version control repository, where there is a record of who submitted what code changes (including Infrastructure as Code changes) and who approved updates. This may also include SAST if static analysis is performed directly off the version control repository, not at the CI/CD server.
  - The CI/CD server, where a range of security tests are performed, and all results should be recorded and issues generated for failed tests.
  - Logs for any artifact repositories, including checked-in builds, testing performed in the repository and repository access.
- Once a deployment is pushed, automated assessment tools—including vulnerability analysis, web application assessment and configuration analysis—should evaluate the live environment. This will likely include use of cloud platform configuration compliance policies, which may offer both detective and preventive controls.

The existence of compliance policies and the quality and completeness or coverage is a meaningful indicator of maturity. The auditor should check for the organization’s policies and standards in this area. The auditor also should check if and how the DevOps teams are dealing with noncompliance and the lead time it takes to remediate issues.

Page intentionally left blank

8.8 Chapter 8 Knowledge Check

REVIEW QUESTIONS

1. Some of the reasons to integrate security in the early stages of the systems development life cycle include (select all that apply):

   A. It can help reduce costs, as security vulnerabilities can be identified before code is deployed in production, reducing the possibility of potential issues and costly fixes.
   B. It helps remediate security vulnerabilities in the infrastructure layers.
   C. It is difficult to understand application requests in order to detect and block attacks outside the application. It is more effective to fix vulnerable code and close off attack vectors.
   D. It helps reduce time for development and deployment of code.

2. What are example benefits of including DAST in software security testing? (Select all that apply.)

   A. DAST can simulate the actions of a malicious user trying to exploit undetected infrastructure vulnerabilities or low patch levels.
   B. DAST can create an open source code inventory and a list of associated vulnerabilities.
   C. DAST has a low rate of false positives.
   D. DAST can simulate how an application reacts to various inputs.

3. Which of the following is not a primary function of a version control repository? (Select the **BEST** answer.)

   A. IAM and approval processes
   B. Source code management
   C. Identifying vulnerabilities and exposures (CVEs) in code
   D. Creating digital fingerprints of known good versions of code

4. What technology should be considered and assessed to manage the extensive credentials typically embedded and used by CI/CD pipelines?

   A. Federated identity brokers to manage user connectivity
   B. Data encryption for the credential files
   C. Secrets managers
   D. Password managers

5. On what elements should the auditor rely for auditability and accountability in a DevSecOps approach? (Select all that apply.)

   A. Major architectural decisions and changes being reflected in the version control system or the issue tracker
   B. Logs of structural changes in the integration pipeline and logs of changes in the build jobs.
   C. Logs of the continuous integration and continuous delivery servers.
   D. Logs of modifications to the compliance registry of the organization.

Answers on page 364

Chapter 8 ANSWER KEY

Correct answers appear in **bold** font.

1. **A. By identifying and remediating code vulnerabilities as early as possible in the development process, the damage of potential exploitation together with impact remediation costs are avoided.**

   B. Remediation of security vulnerabilities in the infrastructure layers should be done in addition to remediating security vulnerabilities in application code. Integrating security in development pipelines does not remediate infrastructure level vulnerabilities.

   C. It is more difficult to build defenses and implement security controls at the infrastructure layers to try protect vulnerable code than to build and deploy secure code.

   D. Integrating security in the development pipeline does not reduce development or deployment time. On the contrary, those times could potentially increase.

2. A. DAST only crawls through the interface and tests various inputs, it does not check for infrastructure vulnerabilities or patch levels. (Section 8.3, Types of integrated security tests)

   B. This is related to source code analysis, not DAST. (Section 8.3, Types of integrated security tests)

   C. (Section 8.3, Types of integrated security tests)

   D. (Section 8.3, Types of integrated security tests)

3. A. Version control tools store code and help with management tasks like versioning, IAM and approval processes (8.4.1.1 Version Control Repository).

   B. Source code management, configuration management databases, container registries and similar type of tools store code and perform management tasks (8.4.1.1 Version Control Repository).

   C. Identifying vulnerabilities is not a function of version control repository. However, several composition analysis vendors integrate their products with source code repository tools to ensure no known CVEs are present in the code (8.4.1.1 Version Control Repository).

   D. Additional facilities to create digital fingerprints for known good versions and other version control features are becoming more common (8.4.1.1 Version Control Repository).

4. A. Federated identity brokers are not generally used for managing stored credentials

   B. While a secrets manager will use encryption internally, data encryption itself is not the right way to protect and manage these credentials.

   C. Secrets managers are specifically designed to manage stored secrets and credentials for applications and automation, like CI/CD pipelines.

   D. CI/CD pipelines store more than just passwords, such as secret access tokens, and password managers are also not suited for automation scenarios.

5. **A. Correct, this is one of the sources of accountability.**

   B. Correct, this is one of the sources of accountability.

   C. Correct, and the auditor should also make sure that detected defects are acted upon.

   D. This is unrelated to DevSecOps.

CCAK™

Certificate of Cloud Auditing Knowledge
A Cloud Security Alliance® and ISACA® Credential

Chapter 9:

Security Trust Assurance and Risk (STAR) Program

Security Trust Assurance and Risk (STAR) Program

Security Trust Assurance and Risk (STAR) Program

9.1 Learning Objectives

After completing this chapter, learners will be able to:

1. Outline the components of the STAR program.
2. Explain the security and privacy implications of STAR.
3. Describe the Open Certification Framework.
4. Recall CSA STAR attestation and certification.
5. Detail STAR continuous auditing.

9.2 Overview

This chapter explains all aspects of the CSA STAR Program. It discusses the building blocks of the STAR Program and how, in the cloud sector, generic programs suffice as a foundation, but sector-specific requirements drive controls that are critical to the users and providers in that sector.

9.3 The STAR Program Components

STAR is a consensus-driven and industry-led program designed to help CSPs and CSCs. It is a governance, risk and compliance program focused on security and privacy that was created to facilitate increasing trust and accountability in the cloud market. In STAR, accountability and trust are defined as functions of increasing levels of security, privacy assurance and transparency.

STAR is a living process that continues to evolve, but it is currently based on the following three pillars, as shown in **figure 9.1**:

- Technical standards: CCM, CSA GDPR Code of Conduct (CoC) for GDPR Compliance
- National and internationally accepted certification and attestation frameworks
- A public registry

Figure 9.1—Three Pillars of STAR Program

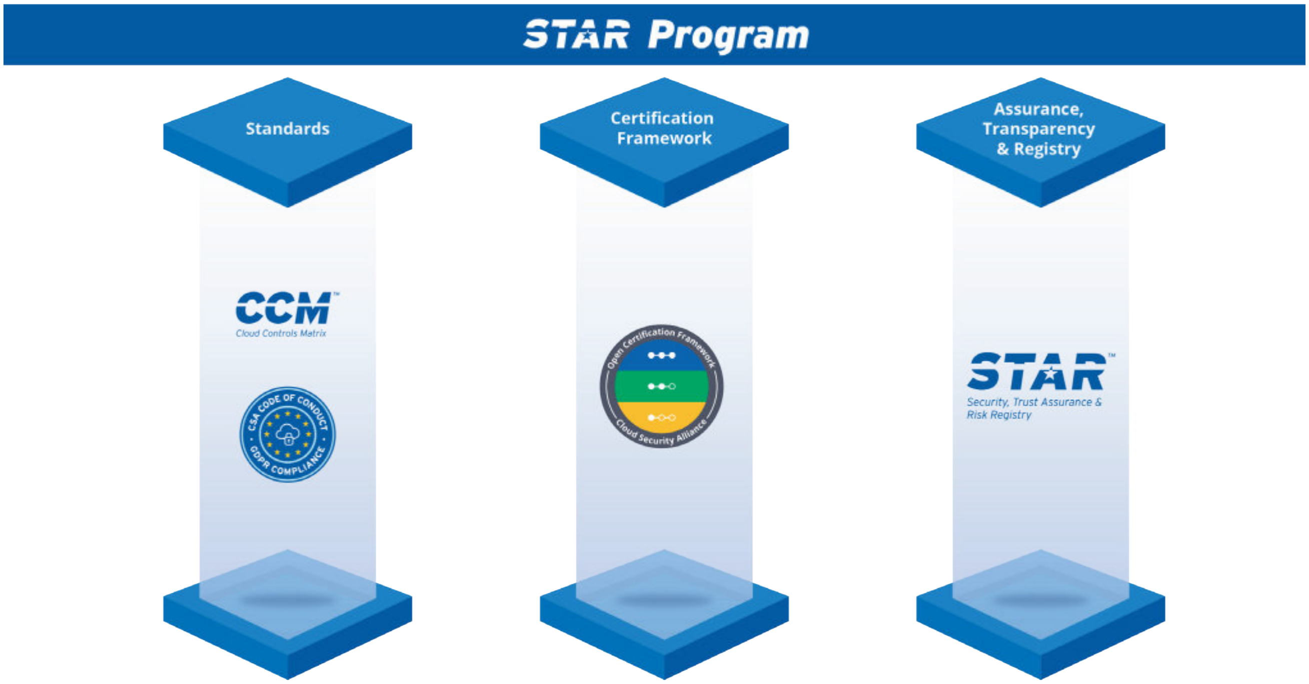
STAR Program
Standards
Certification Framework
Assurance, Transparency & Registry
CCM™
Cloud Controls Matrix
CSA CODE OF CONDUCT GDPR COMPLIANCE
Open Certification Framework
Cloud Security Alliance
STAR™
Security, Trust Assurance & Risk Registry

Source: Cloud Security Alliance, STAR Program

9.3.1 Standards for Security and Privacy

The first pillar of the STAR Program consists of technical best practices for security and privacy developed by the CSA, which have become *de facto* standards over time. The current best practices of reference for the STAR Program are:

- CCM
- CSA Code of Conduct (CoC) for GDPR Compliance

See chapter 3 for details on the CCM. The CSA CoC for GDPR Compliance is described in the following section.

CSA CoC for GDPR Compliance

The CSA CoC for GDPR Compliance (CSA CoC) is best practice for GDPR compliance and a transparency guideline about the level of data protection that the CSP offers relating to its services.

The CSA CoC was developed within CSA by an expert working group (WG) composed of representatives of CSPs, local supervisory authorities, and independent security and privacy professionals (PLA Working Group). It was first published on 21 November 2017, and last updated on 30 September 2020 to reflect the European Data Protection Board's (EDPB) most recent guidelines[260] and to incorporate by reference the *Commission nationale de l'informatique et des libertés* (CNIL) Guides – 2018 Edition: Security of Personal Data.

[260] Former Article 29 Data Protection Working Party

The primary audience for the CSA CoC consists of CSPs rather than cloud customers, because only CSPs can submit their services for adherence to the CoC. The CSA CoC design seeks to benefit CSPs and CSCs (and data subjects and the cloud community in general), by means of the controls imposed through the privacy level agreement (PLA) control specifications. The CSA CoC aims to provide the following:

- For cloud customers of any size, a tool to evaluate the level of personal data protection in services that various CSPs provide (thereby supporting informed decisions about adopting those services)
- For CSPs of any size and geographical location, guidance on complying with European Union (EU) personal data protection legislation and disclosing, in a structured way, the level of personal data protection they offer to customers in connection with their services

The CoC seeks to create additional value for potential and current cloud customers, and for CSPs, data subjects, and the cloud computing community at large, through the following:

- Identifying—in an organic, structured, and systematic manner—all relevant GDPR provisions that CSPs must comply with when handling personal data
- Explaining the GDPR provisions and their practical relevance when applied to the computing environment, while considering clarifications from the European Data Protection Board, and guidance by EU national supervisory authorities
- Emphasizing the need for transparency and enabling compliance with the principle of accountability for CSPs by establishing a disclosure policy and requiring that CSPs adhere to the CoC requirements to provide minimum information and evidence to demonstrate their compliance
- Allowing for public scrutiny of compliance with the CoC, by requiring participating CSPs to disclose their CoC self-attestation/third-party assessment submissions to the CSA STAR Registry

Scope and Methodology

CSA CoC deals only with the business-to-business (B2B) sector and considers cloud customers as companies rather than individuals (as opposed to the business-to-consumer, or B2C market). It addresses specific services a CSP provides in a B2B context—CSPs may offer a variety of services, some of which comply with the CoC terms, and others that do not.

The processing activities related to services covered by the CoC are those a CSP performs when acting as a processor on behalf of a cloud customer. This typically reflects one of two main types of situations:

- The cloud customer acts as a controller, and the CSP acts as a processor on its behalf.
- The cloud customer acts as a processor (on behalf of another party), and the CSP acts as a subprocessor on its behalf.

Regardless of whether the CSP acts as a processor or subprocessor, it must comply with the controls laid out in the CoC to allow the cloud customer (acting as a controller or processor itself) to consider those controls when crafting any offerings that may include the CSP services.

The CSA CoC is based on the mandatory legal provisions of the applicable EU personal data protection framework. It reflects the relevant interpretation by the European supervisory authorities and related best practices developed by relevant agencies. It aims to be a tool for achieving or assessing compliance with the EU personal data protection legislation horizontally across different sectors and domains.

Components

The CoC is structured in three parts:

- Part 1 of the CoC describes its objectives, scope, methodology and assumptions, and provides explanatory notes.

- Part 2 of the CoC contains the control specifications: PLA and its substantial provisions, developed by the PLA WG.
- Part 3 of the CoC outlines the governance structure and the mechanisms of adherence to the CoC.

Qualification as a Draft Code of Conduct Pursuant to Article 40 GDPR

The CoC specifies how to apply the GDPR in the cloud environment, primarily focusing on the following requirements:

- Fair and transparent processing of personal data
- The information provided to the public and to data subjects (as defined in Article 4 (1) GDPR)
- The exercise of the rights of the data subjects
- The measures and procedures referred to in GDPR Articles 24 and 25 and the measures to ensure security of processing referred to in GDPR Article 32
- The notification of personal data breaches to supervisory authorities (as defined in Article 4 (21) of GDPR) and the communication of such personal data breaches to data subjects
- The transfer of personal data to third countries

The CoC—and in particular the Control Specifications: PLA—further seeks to address several key characteristics of the cloud domain regarding privacy and data protection, including specific issues that arise, e.g.:

- Resource-sharing issues, which may trigger potential conflicts due to the use of shared systems and infrastructures to process personal data related to multiple different cloud customers and types of data subjects
- The shared responsibility model, as referenced in previous sections, which sometimes creates confusion between the parties, especially about the difference between transferring the responsibilities for implementing certain security controls to the CSP and maintaining the accountability and duty of care to the cloud customers
- Law enforcement requests, because CSPs located in multiple jurisdictions may find themselves legally required to disclose personal data about cloud customers to public authorities, potentially in breach of EU data protection law
- Complex outsourcing chains that may lead to CSPs engaging a multitude of subprocessors located in different territories worldwide, creating contracting arrangements that may be difficult to manage, and often not providing cloud customers with sufficient means to oppose use of their data in this way (including by not affording sufficient transparency to those cloud customers)
- Obstacles to addressing data subject requests that cloud services may not have been properly configured to allow or comply with promptly and properly, e.g., requests for erasure, restriction of processing, or portability
- Changes to a cloud service's terms, which may affect the way a cloud customer's data are processed in order to provide the service
- Lack of isolation, whereby CSPs (due to their control over data they handle on behalf of multiple cloud customers) could potentially decide to connect data for secondary purposes that its cloud customers may not have authorized
- Further unauthorized processing of data by CSPs in general, such as use of cloud customer data to further develop analytics/profiling algorithms, or for programmatic advertising purposes
- Complexities in allocating data protection roles to cloud customers and CSPs, considering the various types of activities a CSP may carry out (some of which may be as a processor, and others as a controller) and recent Court of Justice of the European Union (CJEU) jurisprudence on joint controllers
- Vendor lock-in, whereby cloud customers may find themselves restricted in their ability to switch between CSPs, e.g., due to a lack of means to ensure portability of personal data between CSP services.

9.3.2 Open Certification Framework

The CSA Open Certification Framework (OCF) is an industry initiative that is the foundation of the CSA STAR Program. The OCF defines and manages the schemes that the STAR Program uses, and allows for global, accredited certification of CSPs. It provides flexible, incremental and multilayered cloud provider certifications and attestations using CSA security guidance and control objectives. The program integrates with national and international third-party certifications and attestations developed within the public accounting community and accredited certifying bodies to avoid duplication of effort and cost.

OCF Structure

The Open Certification Framework is a program based on the CSA standards, security guidance and control objectives, encompassing key principles of transparency, rigorous auditing and harmonization of standards. Companies that use the OCF guidelines demonstrate best practice and validate the security posture of their cloud usage.

The OCF comprises three levels, with the level of transparency and assurance increasing as an organization moves up the ladder of the scale (**figure 9.2**).

Figure 9.2—STAR Levels

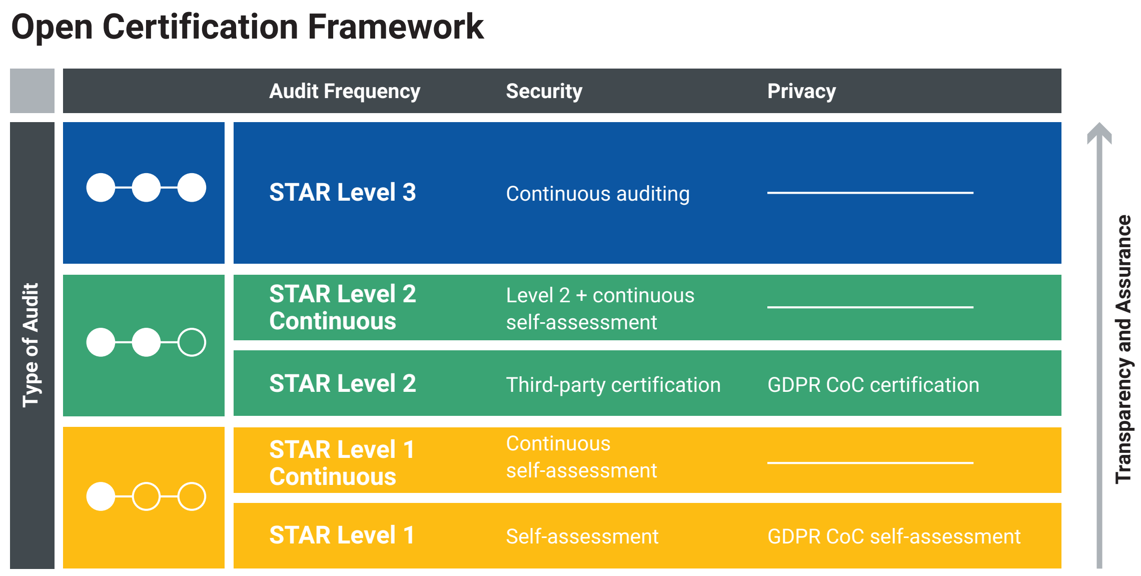

Open Certification Framework

Type of Audit	Audit Frequency	Security	Privacy
●—●—●	STAR Level 3	Continuous auditing	—
●—●—○	STAR Level 2 Continuous	Level 2 + continuous self-assessment	—
	STAR Level 2	Third-party certification	GDPR CoC certification
●—○—○	STAR Level 1 Continuous	Continuous self-assessment	—
	STAR Level 1	Self-assessment	GDPR CoC self-assessment

Transparency and Assurance

Source: Cloud Security Alliance, STAR Program

Figure 9.3 shows the relationship between OCF levels:

- From the assurance perspective, OCF Level 1 provides good-to-moderate assurance, OCF Level 2 provides high assurance, and OCF Level 3 provides very high assurance.
- From a transparency perspective, OCF Level 1 provides good transparency, OCF Level 2 provides low to high transparency, and OCF Level 3 provides very high transparency.

- The degrees of transparency the three OCF levels offer do not necessarily correspond to the three levels of assurance. For instance, OCF Level 1 could provide better transparency than OCF Level 2, because neither the STAR Certification nor STAR Attestation schemes require the organization to make its security controls publicly available.
- CSA encourages organizations aiming to certify at OCF Level 2 to first self-assess at OCF Level 1.

Figure 9.3—Relationships Between OCF Levels

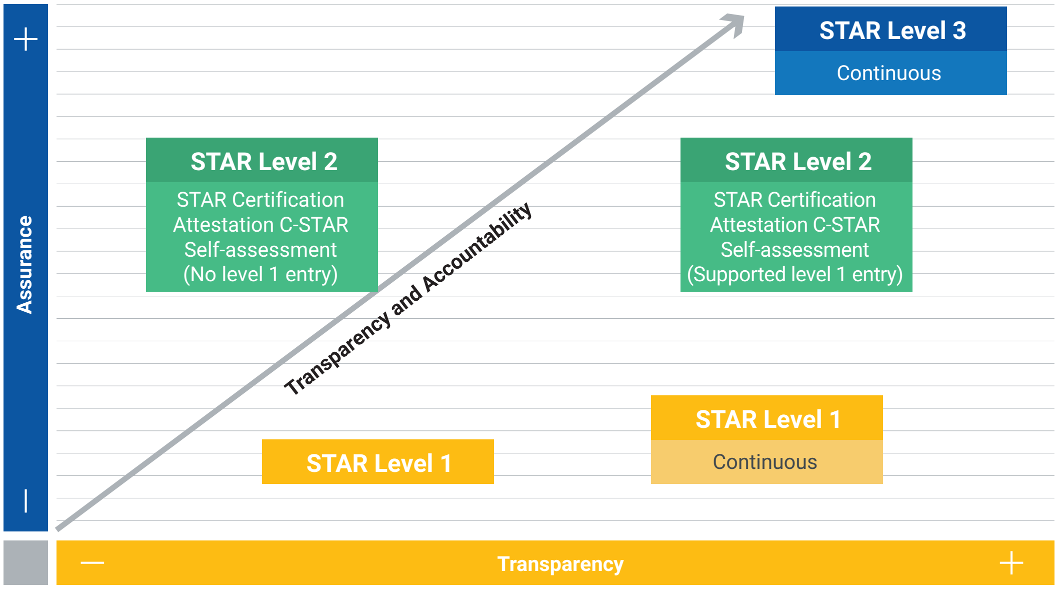

Source: Cloud Security Alliance, STAR Program

Self-Assessment

CSA STAR Self-Assessment is a self-attestation that a CSP carries out to document its security and privacy controls and posture as a cloud service. The attestation is publicly available on the CSA STAR Registry to help customers and potential customers assess the CSP and its cloud services. Via the STAR Self-Assessment, a CSP can document the security controls provided by cloud computing offerings, thereby helping users to assess the security and privacy of CSPs they currently use or are considering using. The self-assessment is based on the Consensus Assessments Initiative Questionnaire (CAIQ)[261] and privacy level agreement (PLA) code of practice (CoP) and serves the purpose of documenting compliance with the Cloud Controls Matrix (CCM)[262] and CSA GDPR Code of Conduct.

Certifications, Attestations and C-STAR

Level 2 of STAR consists of a cloud-relevant third-party audit-based certification and attestation built on CSA best practices and established international auditing standards such as SOC 2[263] and ISO/IEC 27001.

[261] Cloud Security Alliance, "Consensus Assessment Initiative Questionnaire (CAIQ) v3.1," 1 April 2020, https://cloudsecurityalliance.org/artifacts/consensus-assessments-initiative-questionnaire-v3-1/
[262] Cloud Security Alliance, "Cloud Controls Matrix (CCM)," https://cloudsecurityalliance.org/research/cloud-controls-matrix/
[263] SOC 2 is an auditing procedure that ensures service providers securely manage data to protect enterprise interests and client privacy.

- **CSA STAR Attestation**—CSA STAR Attestation is an auditing procedure to report on the examination of the implementation of trust service principles (TSP) and cloud-specific control objectives (CCM). CSA STAR Attestation can be considered as a SOC 2 Type 2 attestation augmented by CCM requirements. It was created thanks to a collaboration between CSA and the American Institute of CPAs (AICPA) to provide guidelines for CPAs to conduct SOC 2 engagements using criteria from the AICPA (Trust Service Principles, AT 101) and the CSA Cloud Controls Matrix.
- **CSA STAR Certification**—The CSA STAR Certification is a third-party independent assessment of a CSP's security using the technology-neutral requirements of the ISO/IEC 27001:2013 management system standard together with the CSA Cloud Controls Matrix. STAR Certification is valid for three years and expires unless updated. In addition, STAR Certification provides a maturity scoring mechanism to improve the assessment of the organization on how well its management system is being managed, and it provides the auditee with an internal report of its strengths and weaknesses.
- **CSA C-STAR Assessment**—The CSA C-STAR Assessment is another third-party independent assessment of the security of a CSP, but it is specifically designed for China-based companies based on China's national standards. C-STAR leverages the requirements of the GB/T 22080-2008 management system standard together with the CSA Cloud Controls Matrix, plus 29 related controls selected from GB/T 22239-2008 and GB/Z 28828-2012. Certificates expire after three years, unless updated. There is no maturity assessment with C-STAR.
- **CSA GDPR COC third-party audit-based certification**—The third-party certification, which is available in 2021, covers the same scope as the self-assessment, but rather than being a self-attestation, a CoC third-party assessment is obtained by having a qualified CoC auditing partner validate a CSP's adherence to the control specifications.

Continuous Assessment, Audit and Certification/Attestation

STAR Level 3 enables automation of the current security practices of a CSP. Providers publish their security practices according to CSA specifications. Customers and tool vendors can retrieve and use this information in a variety of contexts.

STAR Level 3 introduces additional levels of assurance and transparency of the cloud security management system through more frequent testing. These concepts are described as continuous auditing and continuous certification and are discussed in detail in this chapter. STAR Level 3 provides CSPs with a cost-effective way to better communicate the effectiveness of their security program and improve transparency to customers.

9.3.3 STAR Registry

The STAR Registry is a transparency and assurance tool for cloud computing. Launched by CSA in 2011, it has become the trusted and authoritative repository of cloud governance risk- and compliance-related data. The registry consists of a collection of data about the security and privacy posture of cloud services based on CSA best practices as they relate to STAR Level 1 and Level 2 self-assessment and third-party audits. The data in the registry is publicly accessible and searchable, and it allows existing and potential cloud customers to review the security and privacy controls provided by cloud computing offerings (**figure 9.3**).

The main goals of the STAR Registry:

1. To help cloud users gain visibility into the security and privacy capabilities of a wide range of services, allowing them to make informed and risk-based decisions. The information in the registry is structured according to a standard format (CAIQ) allowing users to consistently compare different services to each other.
2. To allow CSPs to prove their compliance commitments in terms of security, privacy and transparency to users and regulators. In addition, for the CSP the registry represents a tool for reducing the effort required to address customers' requests for information (e.g., customer questionnaires, RFIs, RFPs). In essence, rather than

answering each customer priority vendor's questionnaire, CSPs can direct all inquiries to the information published in the STAR Registry.

3. To allow any member of the cloud supply chain (IaaS, PaaS, SaaS) to better understand the allocation of shared responsibilities and build an effective accountability program.

It is prudent that cloud users require CSPs to submit a CAIQ self-assessment to the CSA STAR Registry (**figure 9.4**).

Figure 9.4—STAR Registry

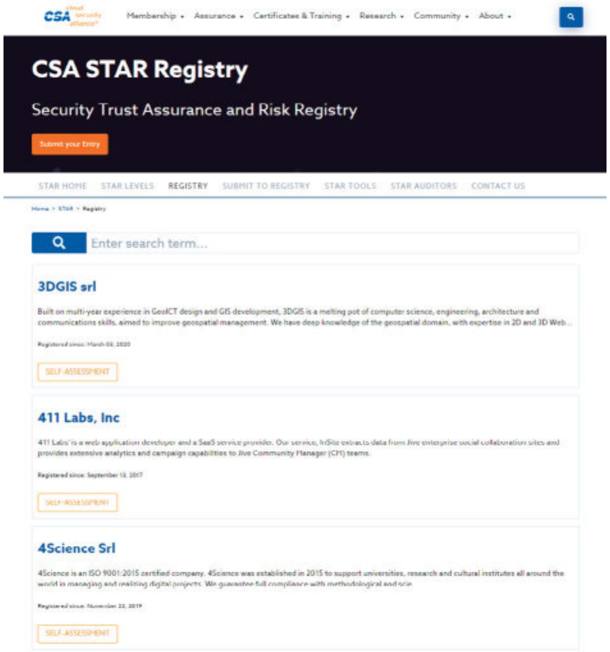

CSA STAR Registry

Security Trust Assurance and Risk Registry

3DGIS srl

Built on multi-year experience in GeoICT design and GIS development, 3DGIS is a melting pot of computer science, engineering, architecture and communications skills, aimed to improve geospatial management. We have deep knowledge of the geospatial domain, with expertise in 2D and 3D Web...

Registered since: March 03, 2020

411 Labs, Inc

411 Labs' is a web application developer and a SaaS service provider. Our service, InSite extracts data from Jive enterprise social collaboration sites and provides extensive analytics and campaign capabilities to Jive Community Manager (CM) teams.

Registered since: September 13, 2017

4Science Srl

4Science is an ISO 9001:2015 certified company. 4Science was established in 2015 to support universities, research and cultural institutes all around the world in managing and realizing digital projects. We guarantee full compliance with methodological and scie

Registered since: November 22, 2019

Source: Cloud Security Alliance, STAR Program

STAR Registry API

CSA has developed an application programming interface (API) that provides access to the content of the STAR Registry in machine-readable format. It offers access to the content of more than 600 CAIQ self-assessments submitted by cloud providers across the globe.

The API allows querying in real time of the following:

- The list of CSPs that have a cloud service in the STAR Registry
- The list of cloud services in the STAR Registry
- Whether a cloud service was subject to a self-assessment, a STAR certification or attestation, with a link to supporting documents
- If the cloud service was subject to a CAIQ self-assessment, the details of the response provided for each of the 295 questions in the CAIQ in machine-readable format

The API provided is based on the REST paradigm and uses JSON, enabling developers to enhance cloud security products, procurement applications and cloud monitoring tools with valuable assurance data.

Access to CAIQ responses in a machine-readable format through the API is possible only if the CSP has provided a CAIQ submission using the standard Excel spreadsheet template provided by CSA or exported from the STARWatch tool. A few CSPs have selected nonportable formats, and their CAIQ assessments answers are therefore not accessible in machine-readable format. CSA encourages providers to submit data in a processable format, because it enables their assurance data to be accessible to the tools that developers build using the CSA API, benefiting their brand.

The API will continue to evolve to support STAR Continuous Auditing and certification as well as STAR Level 3 Continuous Monitoring.

Note:

- The STAR Registry does not provide details of noncompliance, to avoid potentially compromising the security of cloud services under scrutiny.
- A noncompliance does not lead to the immediate removal of a cloud service from the STAR Registry. Cloud providers are offered a grace period that allows them to correct the noncompliance or fix any issues that caused reporting failures, e.g., tool or network issues.

9.4 STAR Level 1

This section provides details about STAR Level 1.

9.4.1 Self-Attestation

Organizations can submit a self-attestation concerning both security and privacy.

Security Self-Attestation

The security self-attestation is the voluntary publication on the CSA STAR Registry of the assessment a CSP has conducted of its cloud services using the CAIQ. As detailed in chapter 3, the CAIQ is an extended question set that customers, third-party assessors and regulators might ask to verify the alignment of the cloud service or CSP with the requirements included in the Cloud Controls Matrix.

The CSP that decides to adhere to STAR Level 1 with the security self-attestation needs to complete the CAIQ questionnaire by answering the 310 questions included in the CAIQ. The questions can be answered as Yes/No/Not Applicable. The CSP may provide further details to justify its answers.

The submission to the security self-attestation is open to any organization and is valid for 12 months. The submission process is via the CSA website, and CSA personnel check the document to ensure it has been completed correctly.

Privacy Self-Attestation

The privacy self-attestation is based on the requirements of the CSA Code of Conduct for the General Data Protection Regulation (GDPR) Compliance described in 9.2.1.1 and carried out using the privacy level agreement submission template.

To adhere to the CoC, the CSP must implement and describe technical and organizational measures that are objectively sufficient to meet every single one of the CoC's controls (Part 2). This allows the CSP to describe the level of privacy and data protection that it undertakes to maintain for the relevant personal data processing activities it performs, regarding the services it provides to the cloud customers the CSP has aligned with this CoC.

The CoC contains 15 areas of requirements, referred to as controls, which are divided into subcontrols. The 15 controls follow:

1. CSP declaration of compliance and accountability
2. CSP relevant contract and its role
3. Ways in which the data will be processed
4. Recordkeeping
5. Data transfer
6. Data security measures
7. Monitoring
8. Personal data breach
9. Data portability, migration and transfer back
10. Restriction of processing
11. Data retention, restitution and deletion
12. Cooperation with the cloud customers
13. Legally required disclosure
14. Remedies for cloud customers
15. CSP insurance policy

A CSP may offer a variety of services to cloud customers. The CoC does not apply to a CSP itself (as an entity), but rather to one or more of the services it offers. Therefore, it is possible for a CSP to fulfill the requirements of this code for many of its services, but still provide other offerings this code does not cover.

The Code of Conduct Self-Assessment is evidence-based and includes voluntary publication on the CSA Security, Trust Assurance and Risk Registry (CSA STAR). It consists of two documents:

- Self-attestation statement of adherence[264]
- Self-attestation results based on the privacy level agreement (PLA) code of practice (CoP)[265]

Requirements of the code (and consequently of the GDPR) are based on a thorough evaluation audit performed by a qualified assessor operating on behalf of the CSA monitoring body. During the evaluation audit, the qualified assessor verifies the correct implementation of all the CoP requirements and the accuracy of information included in the CoP template that the submitter populates.

The CSP must update privacy self-assessments every year.

9.4.2 Self-Assessment as a Transparency and Assurance Tool

The challenge with most security and privacy evaluation templates is that security teams at CSPs often create these documents (security and privacy FAQ, RFP, RFO, RFQ questionnaires) based on their own risk understanding, using a list of specific controls based on their internal experiences, or on a security standard that is not relevant to the cloud. They usually apply the same prescription to every use case and vendor, although it may not be applicable.

[264] Cloud Security Alliance, "PLA Code of Conduct (CoC): Statement of Adherence Self-Assessment," 19 November 2019, https://cloudsecurityalliance.org/artifacts/pla-code-of-conduct-coc-statement-of-adherence-self-assessment/

[265] Cloud Security Alliance, "PLA Code of Practice Template Annex 1 (Updated - March 2020)," 12 March 2020, https://cloudsecurityalliance.org/artifacts/pla-code-of-practice-template-annex-1/

The CAIQ/CCM aligns with more than 35 leading standards and regulations, while the CSA CoC is based on the requirements of GDPR that apply to cloud services, and other relevant legislation, recommendations, guidelines, court cases and best practices from European National Data Protection Authorities, EDPB, EU Parliament, EU Commission, EU Court of Justice cases and ENISA.[266] The self-attestation carried out according to these CSA best practices ensures relevance and comparability between different offerings.

An industry-accepted way of documenting the security and privacy controls that exist in cloud services provides security control transparency and, to some extent, assurance.

This helps cloud customers gauge the security posture of prospective CSPs and determine if their cloud services are suitably secure. It allows the cloud user to review provider security practices, accelerating their due diligence and leading to a higher-quality procurement experience.

The customer can easily monitor the provider's ongoing compliance posture, because it is posted on the STAR public registry and updated on a regular basis. This provides greater peace of mind for the customer.

In addition, although the self-attestations, almost by definition, offer limited assurance on the trustworthiness of the statements, such limitations are mitigated by the full documentation of the security and privacy self-assessment being made publicly available on the STAR Registry, exposing it to public scrutiny, and therefore creating a disincentive to misrepresenting a service or organization's actual security or privacy posture.

9.5 STAR Level 2

This section provides details about STAR Level 2.

9.5.1 Understanding STAR Certification (Third-Party)

The CSA STAR Certification is a rigorous third-party independent assessment of the security of a CSP, leveraging the requirements of the ISO/IEC 27001:2013 management system standard with the CSA Cloud Controls Matrix. In addition, CSA STAR has a maturity model assessment that is internal to the organization. Because not all processes are created equal, this report outlines strengths and weaknesses, allowing an organization to concentrate on improving areas of weakness and exploiting strengths. The levels—bronze, silver and gold—represent how well the process is managed but have no connection to how secure an organization (**figure 9.5**).

[266] Cloud Security Alliance, "Cloud Security Alliance Code of Conduct for GDPR Compliance (Updated - May 2019)," 3 June 2019, https://cloudsecurityalliance.org/artifacts/cloud-security-alliance-code-of-conduct-for-gdpr-compliance/

Figure 9.5—Maturity Dashboard

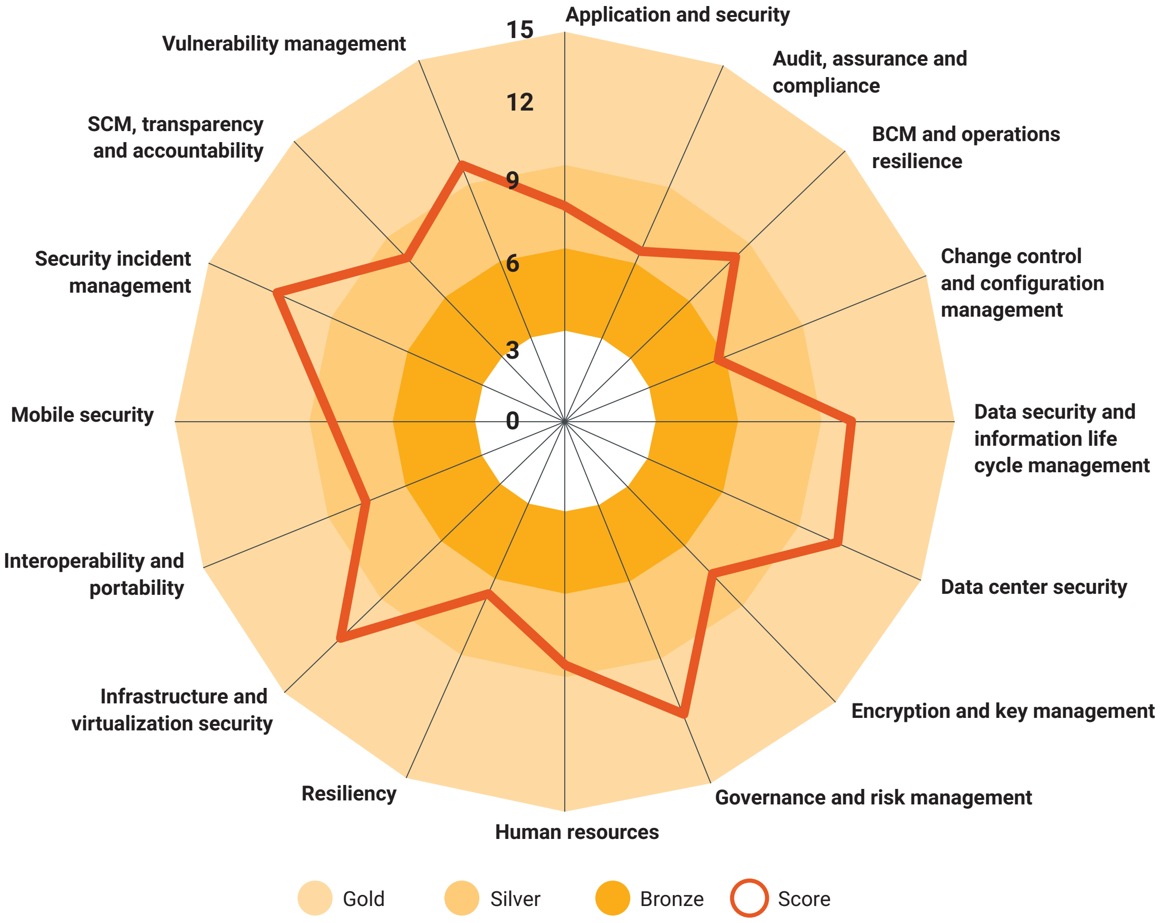

Source: Cloud Security Alliance, STAR Program

All certifying bodies must be accredited to ISO/IEC 17021-1 *Conformity assessment – Requirements for bodies providing audit and certification of management systems* and ISO/IEC 27006 *Information technology – Security techniques – Requirements for bodies providing audit and certification of information security management systems*.

CSA acts as the accreditation body for CSA STAR Certification and has published two scheme documents: Requirements for Bodies Providing Star Certification and STAR Certification Guidance Document: Auditing the Cloud Controls Matrix (CCM).

9.5.2 Purpose of STAR Certification and the Associated Management Capability Model

CSA Star Certification gives a prospective customer a greater understanding of the level of control in place at a certified organization it is considering as a CSP.

It also highlights areas where an organization might wish to improve.

It ensures that the CCM does not become the minimum requirement, but through the model highlights what best-in-class performance is like.

There are both internal (business improvement) and external (customer reassurance and transparency) reasons for auditing to a management capability model.

One of the key objectives of the scheme is to ensure the scope of the CSP is fit for purpose and the process is driven by SLAs and therefore is customer-focused.

When an organization is audited, a management capability score (**figure 9.6**) is assigned to each control area in the CCM, indicating the capability of the management to ensure that the control is operating effectively.

The management capability of the controls is scored on a scale of 1-15. These scores are divided into five categories that describe the type of approach characteristic of each group of scores (see **figure 9.6**).

Figure 9.6—Maturity Scoring Grid

Score	1 to 3	4 to 6	7 to 9	10 to 12	12 to 15
	No Formal Approach	Reactive	Proactive	Improving	Optimizing
Evidence/ Definition	There is no evidence of a system in place to manage the control area.	There is evidence of a system in place to cover operations in the control areas. Where required, the system is documented.	There is evidence of a robust system in place that covers all routine operations in the control area.	There is evidence the system for managing the control area is capable of managing contingency events as well a routine activity.	Control area owners can demonstrate that they actively review best practice from their industry and across their organization and apply it to the control area.
Managed	There is some evidence of either a documented system or an accepted way of working is in place.	There is a clearly identified owner of the control area who understands their scope of responsibility.	There is evidence that the control area is actively monitored and measured and action evaluated based on the evidence.	Input from a variety of sources is considered to decide how and to manage risk and improve operations in this control area.	Control area owners actively share best practice to support development in other areas of the organization based on their experience in this control area.
Followed/ Effective	There is some evidence of an accepted way of working that is broadly understood and followed.	There is evidence the system is understood and routinely followed.	There is evidence that critical people operating in the control area are appropriately trained/skilled to manage routine operations in the control area.	There is evidence that inputs from a range of stakeholders and monitoring and measuring systems have been taken into account when improving operations in the control area.	Changed in the control area are evaluated against strategic objective of the organization.

Source: Cloud Security Alliance, STAR Program

9.5.3 Understanding STAR Attestation (Third-Party)

CSA STAR Attestation is the first cloud-specific attestation program designed to meet this sector-specific need. CSA STAR Attestation is a collaboration between CSA and the AICPA to provide guidelines for CPAs to conduct SOC 2 engagements using criteria from the AICPA (Trust Service Principles, AT 101) and the CSA Cloud Controls Matrix.

Star Attestation is a SOC 2 engagement with the following criteria:

- The applicable criteria in TSP section 100, Trust Services Principles, Criteria, and Illustrations for Security, Availability, Processing Integrity, Confidentiality, and Privacy (AICPA, Technical Practice Aids) (TSPC)
- The control specifications included in the CSA CCM. (The CCM control specifications constitute suitable criteria, as defined by AT Section 101 Attest Engagements (AICPA professional standards) and are referred to in this document as the CCM criteria. AT section 101 is included in the AICPA Statements on Standards for Attestation Engagements, commonly referred to as the attestation standards).

Based on the SOC 2 attestations standards or international equivalents (i.e., ISAE 3000), the STAR Attestation provides a mechanism for reporting on the CSP description of its systems and controls, including a description of the service auditor's tests of controls. The reports are intended to meet the needs of a broad range of users that need to understand the cloud-specific controls of a CSP as they relate to security and the criteria in CCM. As with a traditional SOC 2 attestation, there are two types of reports:

- **Type 1**—Report on management's description of a CSP system and the suitability of the design of controls (point-in-time assessment)
- **Type 2**—Report on management's description of a CSP system and the suitability of the design and operating effectiveness of controls (over-a-period-of-time assessment)

The Type 1 audit is used as a steppingstone to the more rigorous Type 2 audit. A CSA STAR Attestation Type 1 status demonstrates to the CSP's customers the commitment to cloud security through a thorough assessment of the policies and procedures in place to protect the confidentiality, integrity and availability of their data.

A STAR Attestation obtained based on a SOC 2 Type 1 report is valid for six months from the as-of date, i.e., an organization that received its STAR Attestation based on a SOC 2 Type 1 report is required to submit a SOC 2 Type 2 report to maintain uninterrupted STAR Attestation status. The validity period of a STAR Attestation is one year and may be extended by a grace period of three months beyond the basic validity period for report generation and delivery (maximum validity period).

- The CCM control specifications constitute suitable criteria as defined by the attestation standard (CCM criteria) and include criteria equivalent to the criteria for the security principle in the TSC, plus certain additional criteria related to security.
- In a SOC 2 report, the CPA expresses an opinion on whether the controls were operating effectively to provide reasonable assurance that the cloud provider's service commitments and system requirements were achieved based on the applicable trust services criteria and CCM criteria.
- STAR Attestation facilitates reducing complexity in terms of time, cost, trust and transparency.

All assessment firms must follow the CSA STAR Attestation Guidelines: Guidelines for CPAs Providing CSA STAR Attestation v2.

9.5.4 Understanding C-STAR Assessment (Third-Party)

The CSA C-STAR Assessment is a third-party, independent assessment of the security of a CSP for the Greater China market that harmonizes CSA best practices with Chinese national standards (**figure 9.7**). Although STAR Certification is the predominant certification in the APAC region, this is typically for organizations based in or doing a significant amount of business in China that may have a mandate from the Chinese government or a customer. C-STAR leverages the requirements of the GB/T 22080-2008 management system standard with the CSA Cloud Controls Matrix, plus 29 related controls selected from GB/T 22239-2008 and GB/Z 28828-2012 (**figure 9.8**). Certification certificates are valid for three years unless updated.

Figure 9.7—Inputs to C-STAR Assessment Scheme

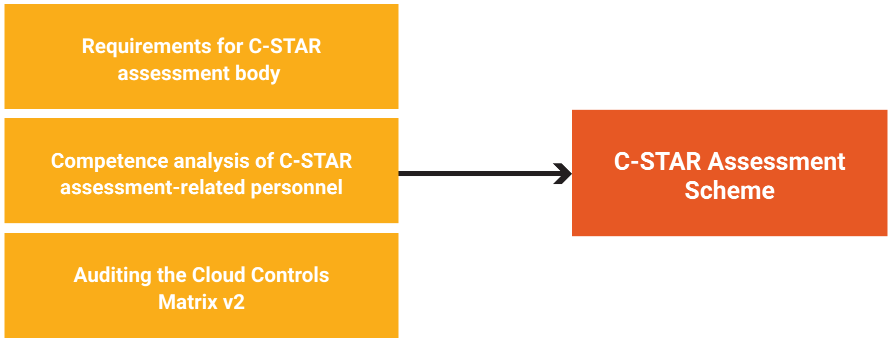

Requirements for C-STAR assessment body

Competence analysis of C-STAR assessment-related personnel

Auditing the Cloud Controls Matrix v2

C-STAR Assessment Scheme

Figure 9.8—C-Star Scheme

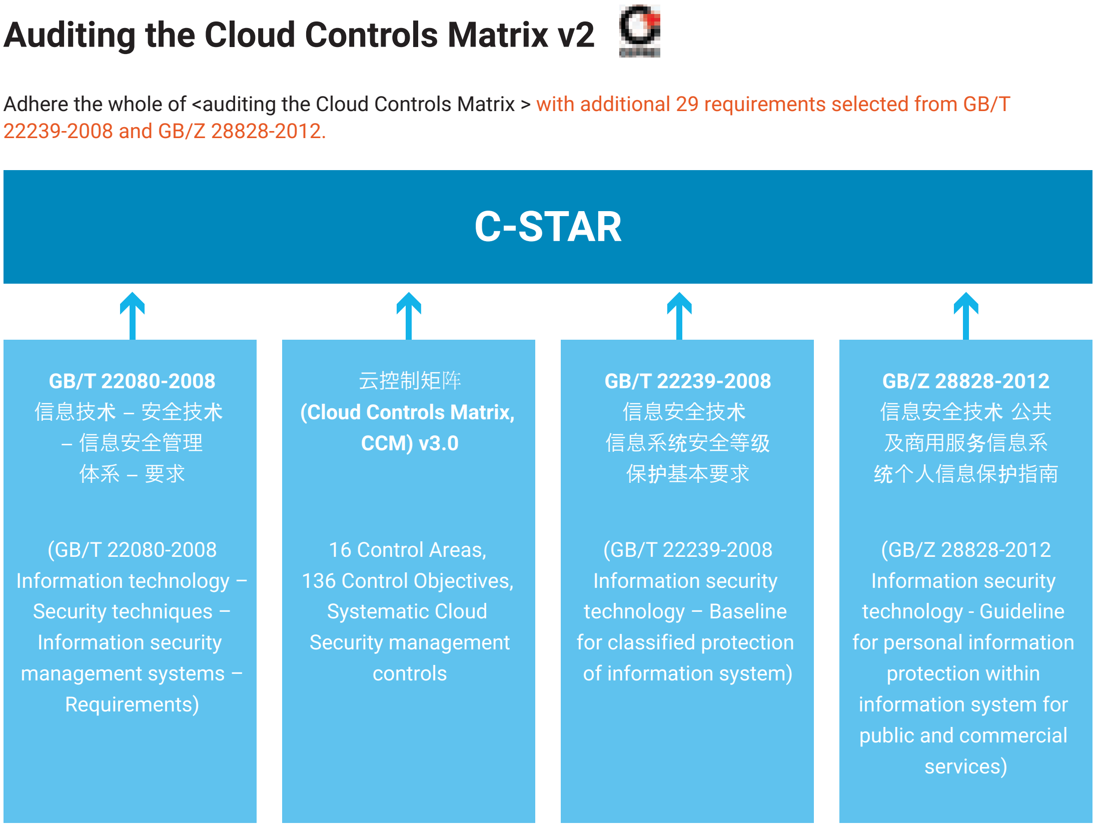

Auditing the Cloud Controls Matrix v2

Adhere the whole of <auditing the Cloud Controls Matrix > with additional 29 requirements selected from GB/T 22239-2008 and GB/Z 28828-2012.

C-STAR

GB/T 22080-2008
信息技术 – 安全技术 – 信息安全管理体系 – 要求

(GB/T 22080-2008 Information technology – Security techniques – Information security management systems – Requirements)

云控制矩阵
(Cloud Controls Matrix, CCM) v3.0

16 Control Areas, 136 Control Objectives, Systematic Cloud Security management controls

GB/T 22239-2008
信息安全技术 信息系统安全等级保护基本要求

(GB/T 22239-2008 Information security technology – Baseline for classified protection of information system)

GB/Z 28828-2012
信息安全技术 公共及商用服务信息系统个人信息保护指南

(GB/Z 28828-2012 Information security technology - Guideline for personal information protection within information system for public and commercial services)

Source: Cloud Security Alliance, STAR Program

Bodies providing C-STAR assessment services must be accredited by the China National Accreditation Service (CNAS) to perform GB/T 22080 certification and meet the capability and reliability requirements.

C-STAR follows a similar certification mechanism as CSA STAR Certification and includes a maturity model, but it follows a certification and assessment guide written specifically for and unique to C-STAR (**figure 9.9**).

Figure 9.9—Comparison Between STAR and C-STAR Schemes

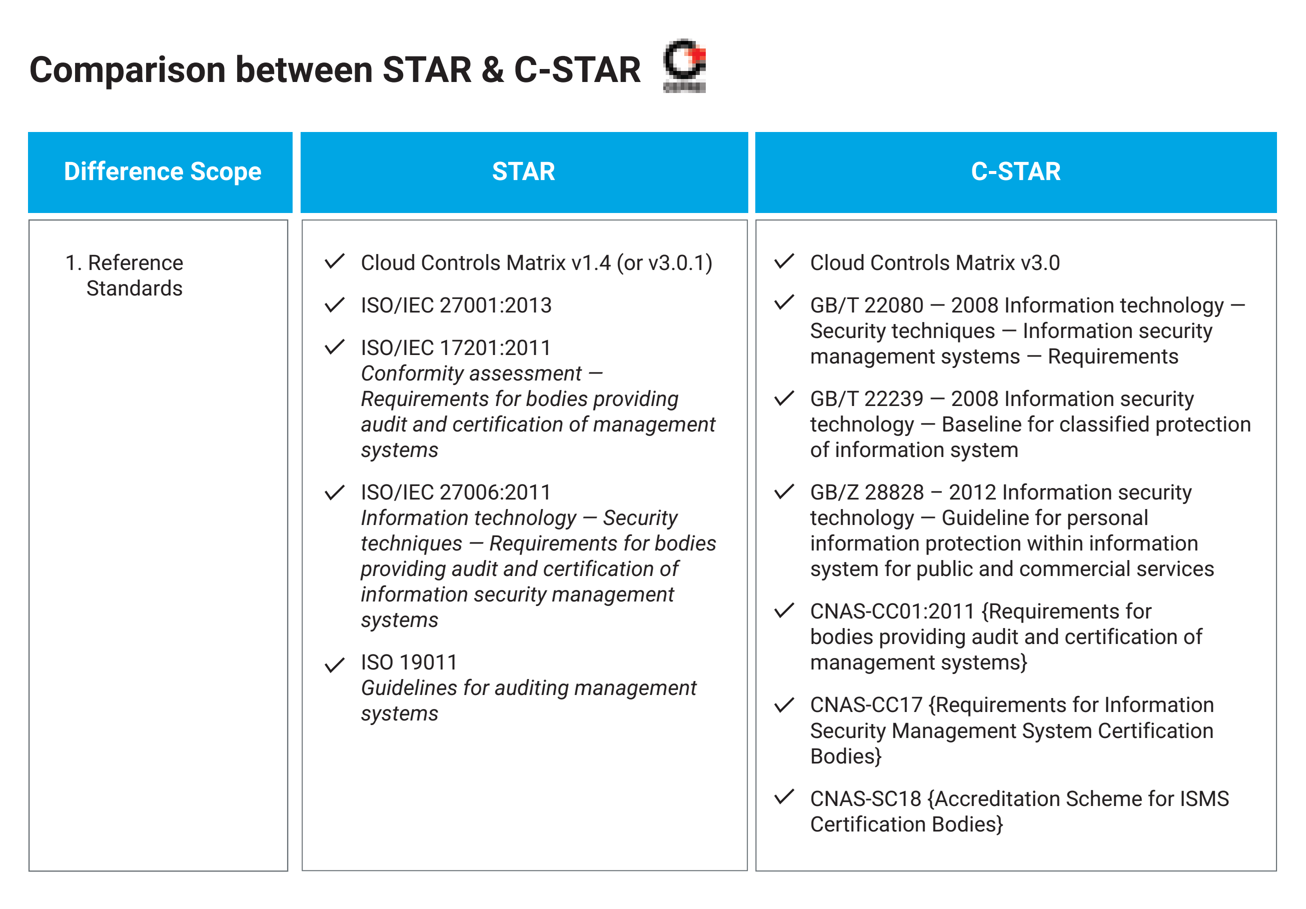

Comparison between STAR & C-STAR

Difference Scope	STAR	C-STAR
1. Reference Standards	✓ Cloud Controls Matrix v1.4 (or v3.0.1) ✓ ISO/IEC 27001:2013 ✓ ISO/IEC 17201:2011 *Conformity assessment — Requirements for bodies providing audit and certification of management systems* ✓ ISO/IEC 27006:2011 *Information technology — Security techniques — Requirements for bodies providing audit and certification of information security management systems* ✓ ISO 19011 *Guidelines for auditing management systems*	✓ Cloud Controls Matrix v3.0 ✓ GB/T 22080 — 2008 Information technology — Security techniques — Information security management systems — Requirements ✓ GB/T 22239 — 2008 Information security technology — Baseline for classified protection of information system ✓ GB/Z 28828 – 2012 Information security technology — Guideline for personal information protection within information system for public and commercial services ✓ CNAS-CC01:2011 {Requirements for bodies providing audit and certification of management systems} ✓ CNAS-CC17 {Requirements for Information Security Management System Certification Bodies} ✓ CNAS-SC18 {Accreditation Scheme for ISMS Certification Bodies}

Source: Cloud Security Alliance, STAR Program

9.5.5 Understanding CSA GDPR COC Certification (Third-Party)

This is a third-party certification scheme that results in a data protection certification evidenced by a data protection certificate, seals and marks. It is available in 2021.

This scheme is offered globally. The CSA PLA Code of Practice (CoP) is a cloud sector guidance to enable organizations to put in place, as part of the overall information governance infrastructure, a personal information system for maintaining and improving compliance with General Data Protection Requirements and good practice. This is achieved by having a system for the design and implementation of the organization's personal information governance system through the use of processes to meet the organization's own requirements, relevant requirements of interested parties (including but not limited to data subjects, data controllers, data processors and supervisory authorities) and the requirements of CSA PLA collectively. The scheme will be aligned with ISO/IEC 17065.

9.5.6 Recognizing Audit Scope

Scope is critical in the STAR and audit process. In the cloud sector, scope must be fit for purpose, which means customer-focused and SLA-driven. A scope is a description of the services and processes covered by the audit, and it ultimately determines the range of activities and the period (months or years) of records that are to be subjected to an audit examination. Scope also defines what sites are included for a multi-site/multi-service organization (**figure 9.10**):

- The scope of an organization must clearly define the functions that are included in the certification or attestation.
- A misleading scope that excludes an area of an organization that might be assumed to be covered in the scope of registration, for instance, shall not be allowed.
- The scope of SOC 2 or ISO/IEC 27001 must not be less than that of the scope of the STAR Level 2 certification/attestation.
- The scope must be written to reflect as closely as possible the full chain of critical activities that have implications for a customer's data or the service it receives. It will cover the core SLAs that the organization has with its clients.

Figure 9.10—Scope Model

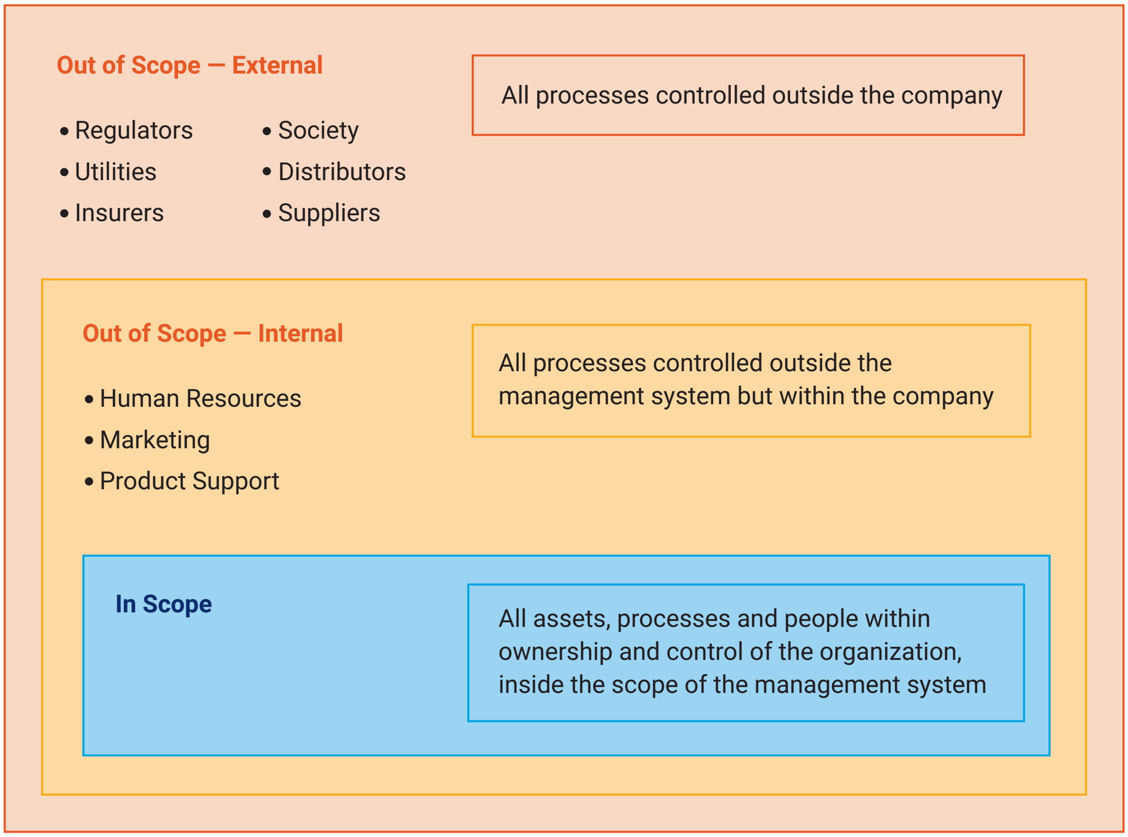
Out of Scope — External
• Regulators
• Utilities
• Insurers
• Society
• Distributors
• Suppliers
All processes controlled outside the company
Out of Scope — Internal
• Human Resources
• Marketing
• Product Support
All processes controlled outside the management system but within the company
In Scope
All assets, processes and people within ownership and control of the organization, inside the scope of the management system

9.6 STAR Level 3

This section provides details about STAR Level 3.

9.6.1 Understanding STAR Continuous

For some cloud customers in sensitive or highly regulated industries, such as banking or healthcare, traditional annual or biannual audits do not provide enough assurance to move to the cloud. To address the concerns of this segment of the industry, the CSA is building STAR Continuous, a framework designed to provide assurance to customers on a monthly, daily or even hourly basis (**figure 9.11**).

Figure 9.11—Continuous Certification

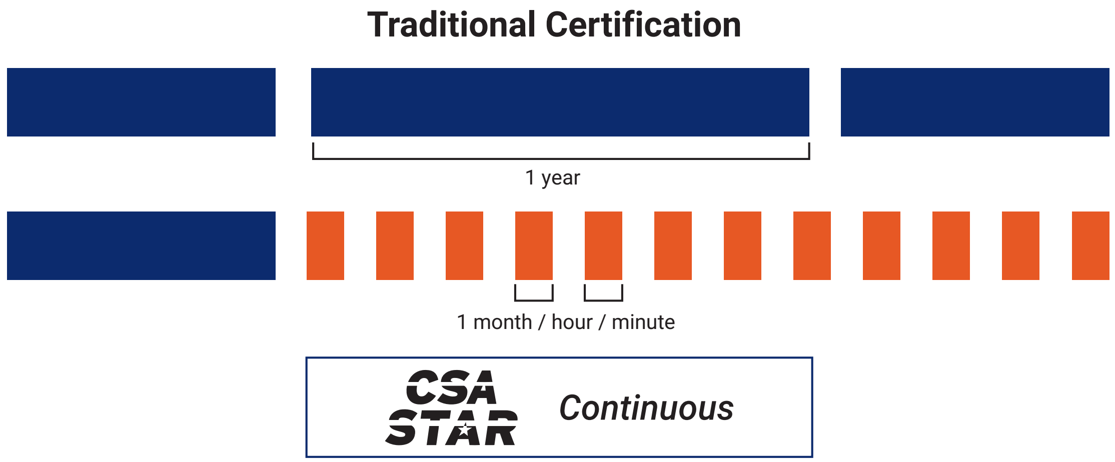

Source: Cloud Security Alliance, STAR Program

The foundation of STAR Continuous is continuous auditing: the ongoing evaluation of certain characteristics of an information system, mostly by automated means, to get near real-time assurance.

STAR Level 3 embodies the highest level of assurance of STAR Continuous by offering a continuous certification of a cloud service, under the supervision of a trusted third party. STAR Level 2 builds upon a point-in-time certification or attestation as provided by STAR Level 2 and extends it with a continuous audit supervised by a third-party auditor.

STAR Level 3 is still in development, and the first certifications are not expected before 2021.

9.6.2 Continuous Auditing Concepts

STAR Continuous is based on a continuous auditing approach being developed by CSA through several EU-funded research projects[267] with partners in the industry, academia and public sector. CSA continuous auditing is built around a few core concepts borrowed from ISO/IEC 19086-1, ISO/IEC 17788 and ISO/IEC 27000:

- **Attribute**—Property or characteristic of an object that can be distinguished quantitatively or qualitatively by human or automated means.
- **Metric**—Standard of measurement that defines the conditions and rules for performing the measurement and for understanding the results of a measurement.

[267] EU-SEC, "The European Security Certification Framework (EU-SEC)," www.sec-cert.eu/

- **Cloud service level objective**—Commitment a cloud service provider (ISO/IEC 17788:2014, 3.2.15) makes for a specific, quantitative characteristic of a cloud service (ISO/IEC 17788:2014, 3.2.8), where the value follows the interval scale or ratio scale.
- **Cloud service qualitative objective**—Commitment a CSP (ISO/IEC 17788:2014, 3.2.15) makes for a specific, qualitative characteristic of a cloud service (ISO/IEC 17788:2014, 3.2.8), where the value follows the nominal scale or ordinal scale.

A continuous audit can be summarized as follows:

- Select attributes of the information system that can be used to evaluate the security of the information system under scrutiny.
- Use metrics to evaluate the attributes and obtain measurement results, based on evidence collected through the information systems through various tools**.**
- Compare the measurement results with predefined objectives, i.e., SLOs and SQOs. Compliant systems are those that meet all SLOs and SQOs.

Ideally, all these actions must be fully automated to make continuous auditing practical and cost-effective. These steps need to be repeated regularly at predefined intervals to provide a continuous relevant picture of the security of the information system being evaluated. However, not all attributes of an information system need to be evaluated at the same frequency. The frequency of evaluation of each security attribute will depend largely on the related risks covered and on technical feasibility. For example, checking that database fields are encrypted might be an activity that must be conducted every few minutes when dealing with very sensitive data. However, checking that disaster recovery procedure documentation is up-to-date is an activity that likely only needs to be performed every few months, not only because changes are unlikely, but also because this is largely a manual check that cannot be automated, contrary to the verification of a technical characteristic such as encryption.

Continuous auditing used to build a continuous certification framework includes the concept of a certification target, a statement comprised of the following:

- A description of the scope of the information under scrutiny
- A list of security attributes being evaluated, each with:
  - A corresponding SLO or SQO
  - A frequency of evaluation

When an organization engages in a continuous certification process, it starts by providing a certification target, which acts as a commitment stating which SLOs and SQOs it intends to comply with, and at what frequency compliance with each SLO or SQO will be assessed and reported (**figure 9.12**).

Figure 9.12—Traditional vs. Continuous Certification: A Conceptual Model Change

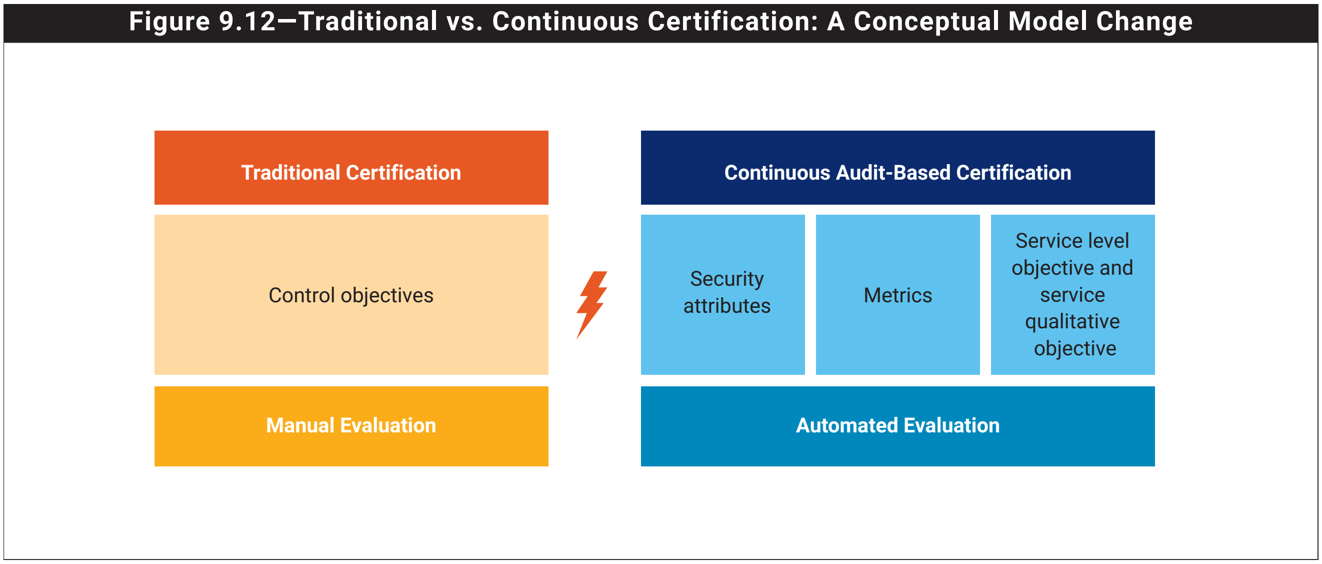
Traditional Certification
Control objectives
Manual Evaluation
Continuous Audit-Based Certification
Security attributes
Metrics
Service level objective and service qualitative objective
Automated Evaluation

9.6.3 Certification Process Overview

The continuous certification scheme for STAR Level 3 extends a traditional certification scheme with a continuous process of automated checks. The whole process can be summarized in two consecutive phases: an initialization phase and a continuous audit phase.

The initialization phase includes:

- The CSP undergoes a traditional third-party audit to obtain a certification or attestation, as detailed for STAR Level 2.
- In addition, the CSP defines:
  - A continuous certification target that comprises a set of security objectives, each associated with a policy that defines the assessment frequency (e.g., check every four hours)
  - A set of tools capable of verifying that the security objectives are fulfilled
- Moreover, the third-party auditor also checks:
  - That the defined continuous certification target covers a satisfactory scope of the certified information system
  - That the reporting tools are trustworthy and fit for purpose
- If this process is successful, the continuous certification target is transmitted to CSA, which acts as a certification authority. CSA creates a corresponding entry for the cloud service in a dedicated section of the STAR Registry that features continuously certified cloud services.

Once the Initialization phase is completed, the continuous audit phase starts, running for up to a full year.

The continuous audit phase includes:

- The third-party auditor periodically performs checks to confirm that the assessment tools are trustworthy (e.g., integrity checks).
- The assessment tools continuously report the results of the assessment of each defined security objective to CSA through a dedicated API, according to the frequency defined in policies within the continuous certification target:

- If a CSP reports in due time that all security objectives are met, the cloud service is marked as “compliant” in the corresponding entry in the public registry.
- If a CSP reports noncompliances or if the CSP fails to report security objectives in due time, the entry will ultimately be removed from the public registry if the situation is not resolved within a predetermined period of time.

Note: Cloud services that are removed from the list of continuously certified service providers due to noncompliance do not lose their point-in-time certification or attestation, which is obtained during the initialization phase.

In the increasingly complex and resource-demanding compliance landscape of the cloud, companies can benefit from sector-specific guidance, tools and controls to manage and optimize their cybersecurity and privacy compliance posture and trusted brand perception. STAR enables CSPs to validate their cloud security and offers proof of the controls they have in place to current and future customers. In turn, STAR helps customers assess which CSPs meet the level of assurance they require.

9.7 Chapter 9 Knowledge Check

REVIEW QUESTIONS

1. What are the **MAIN** benefits of the CSA STAR Registry? (Select all that apply.)

   A. The CSA STAR registry notifies regulatory bodies of the level of compliance of a cloud service with the applicable national standards or regulation.
   B. The CSA STAR registry provides cloud service customers an understanding of the security and privacy posture of a cloud service.
   C. The CSA STAR registry provides a detailed understanding of the shared security responsibilities applicable to cloud service customers and cloud service providers in different cloud technologies.
   D. The CSA STAR registry helps build an effective cloud accountability program.

2. If a CSP would like to perform a self-assessment to understand/show its level of adherence to GDPR, which of the following documents is the **BEST** for the task?

   A. Consensus Assessments Initiative Questionnaire (CAIQ)
   B. Privacy Level Agreement Code of Practice
   C. Cloud Control Matrix (CCM)
   D. International Standards Organization (ISO) 27001

3. CSA continuous auditing is built around four core concepts borrowed from ISO/IEC 19086-1, ISO/IEC 17788 and ISO/IEC 27000. What concepts are included? (Select all that apply.)

   A. Attributes
   B. Metrics
   C. Cloud service level objective
   D. Cloud service qualitative objective

4. The management capability of the controls will be scored on a scale of 1-15. These scores have been divided into five different categories that describe the type of approach characteristic of each group of scores. Name the five categories:

   1. ____________
   2. ____________
   3. ____________
   4. ____________
   5. ____________

Answers on page 388

Chapter 9 ANSWER KEY

Correct answers appear in **bold** font.

1. A. While the CCM mapping feature might provide an understanding of the level of adherence to a certain standard or regulation, there's no notification done to the authority.

   B. The Registry helps cloud users to gain visibility of the security and privacy capabilities of a wide range of services, allowing them to make informed and risk-based decisions.

   C. While the CCM and the STAR Program allow any member of the cloud supply chain to better understand the allocation of shared responsibilities, the info provided do not offer any detailed insight on the specificity of certain technology.

   D. The STAR Program helps members of the cloud supply chain to gain visibility of the security and privacy capabilities of a wide range of services, and it allows to better understand the allocation of shared responsibilities.

2. A. CAIQ is the questionnaire associated with CCM. It covers cloud relevant security and privacy controls, but it is not specifically targeting General Data Protection Regulation (GDPR) requirements.

   B. CoP is the set of privacy controls created for companies to show to third parties and/or to understand its level of preparedness to meet the requirements of General Data Protection Regulation (GDPR) based on the requirements of the CSA Code of Conduct for GDPR Compliance

   C. CCM is the CSA's security control framework for cloud security. It covers certain aspects of privacy, but it is not specifically targeting General Data Protection Regulation (GDPR) requirements

   D. ISO 27001 is the ISO standard defining the requirements on an information security management system, but is neither privacy, nor General Data Protection Regulation (GDPR) specific.

3. **A. Correct.**

   B. Correct.

   C. Correct.

   D. Correct. STAR Continuous is based on continuous auditing, an approach that CSA has developed throughout the years, notably through several EU-funded research projects with partners in the industry, academia and the public sector. CSA's continuous auditing is built around a few core concepts borrowed from ISO/IEC 19086-1, ISO/IEC 17788 and ISO/IEC 27000:

   - **Attribute—Property or characteristic of an object that can be distinguished quantitatively or qualitatively by human or automated means**
   - **Metric—Standard of measurement that defines the conditions and the rules for performing the measurement and for understanding the results of a measurement**
   - **Cloud service level objective—Commitment a cloud service provider (ISO/IEC 17788:2014, 3.2.15) makes for a specific, quantitative characteristic of a cloud service (ISO/IEC 17788:2014, 3.2.8), where the value follows the interval scale or ratio scale.**
   - **Cloud service qualitative objective—Commitment a CSP (ISO/IEC 17788:2014, 3.2.15) makes for a specific, qualitative characteristic of a cloud service (ISO/IEC 17788:2014, 3.2.8), where the value follows the nominal scale or ordinal scale.**

4. CSA's "Auditing the CCM" outlines the maturity model audit and scores that are specified as:

   1. No Formal Approach

   2. Reactive

   3. Proactive

   4. Improving

   5. Optimizing

   See section 9.5.2, "Purpose of STAR Certification and the Associated Management Capability Model," **figure 9.6**.

Page intentionally left blank

Acronyms

The following is a list of common acronyms used throughout the *Certificate of Cloud Auditing Knowledge Study Guide*. These may be defined in the text for clarity.

AES	Advanced Encryption Standard (2.7.1)
AICPA	American Institute of Certified Public Accountants (2.6.2)
ANSSI	French National Cybersecurity Agency (1.4.10)
API	application programming interface (1.4.13)
APP	Australian Privacy Principles (2.6.1)
APPI	Act on the Protection of Personal Information (Japan) (6.4.4)
AS	application service (2.7.6)
ASB	Auditing Standards Board (5.4.1)
AWS	Amazon Web Services (1.4.11)
BCP/DR	Business Continuity Plan/Disaster Recovery (3.4.4)
BCR	Business Continuity Management & Operational Resilience (7.6)
BOSS	Business Operation Support Services (2.7.6)
BSI	Federal Office for Information Security (Germany) (2.6.2)
BSI C5	Federal Office for Information Security (Germany) Cloud Computing Compliance Criteria Catalogue (2.10.2)
BYOK	bring your own key (6.4.4)
C5	Cloud Computing Compliance Controls Catalogue (2.6.2)
CA	Certification Authority (1.4.1)
CA	chartered accountant (5.4.2)
CAIQ	Consensus Assessment Initiative Questionnaire (1.4.10)
CAP	corrective action plan (5.3.1)
CASB	Cloud Access Security Broker (2.4.6)
CCM	Cloud Controls Matrix (2.6.2)
CCOE	cloud center of excellence (1.3.5)
CCPA	California Consumer Privacy Act (1.3.9)
CCSK	Certificate of Cloud Security Knowledge (3.3.4)
CCSP	Certified Cloud Security Professional (5.4.2)
CE	consuming entity (5.5.2)
CFE	Certified Fraud Examiner (5.4.2)
CI/CD	continuous integration/delivery (8.3.5)
CIA	Certified Internal Auditor (5.4.2)
CIO	chief information officer (1.4.12)
CIP	Certified Infrastructure Protection (1.4.10)
CIS	Center for Internet Security (2.6.2)
CISA	Certified Information Systems Auditor (5.4.2)
CISM	Certified Information Security Manager (1.4.12)
CISO	chief information security officer (2.10.3)
CLOUD Act	Clarifying Lawful Overseas Use of Data Act (2.6.1)
COBIT	Controls Objectives for Information and Related Technology (1.3.1)
COPPA	Children’s Online Privacy Protection Rule (3.7.1)
CPA	Certified Public Accountant (5.4.2)
CRM	customer relationship management (1.4.10)
CSA	Cloud Security Alliance (1.3.4)
CSA STAR	Cloud Security Alliance Security Trust Assurance and Risk (1.3.4)
CSC	cloud service customer (2.2))
CSIRT	computer security incident response team (4.5.2)
CSP	cloud service provider (Introduction)
CTO	chief technology officer (2.4.1)
CVE	common vulnerabilities and exposures (8.51)
DAST	dynamic application security testing (8.4.1)
DDoS	distributed denial of service (1.4.14)
Dev	Development (6.6.2)
DMTF	Distributed Management Task Force (2.6.2)
DR	disaster recovery (1.4.9)
DSP	Data Security & Privacy Lifecycle Management (3.8.1)
EDI	electronic data interchange (5.8)
ENISA	European Union Cybersecurity Agency (1.4.10)
ERM	enterprise risk management (1.4.10)
ERMF	enterprise risk management framework (1.4.10)
ETSI	European Telecommunications Standards Institute (Europe) (2.6.2)
EU	European Union (1.4.9)
EU-SEC	European Security Certification Framework (2.6.2)
FaaS	Function as a Service (2.8.2)

FAIR	Factor Analysis of Information Risk (1.4.10)
FedRAMP	Federal Risk and Authorization Management Program (2.6)
FERPA	Family Educational Rights and Privacy Act (2.6.1)
FIDO Alliance	Fast Identity Online Alliance (2.6.2)
FISMA	Federal Information Security Management Act (2.6.1)
FOIA	Freedom of Information Act (1.4.3)
G-Cloud	government cloud (2.8.3)
GAPP	Generally Accepted Privacy Principles (3.7.1)
GCP	Google Cloud Platform (2.7.6)
GCSA	GIAC Cloud Security Automation (5.4.2)
GDPR	General Data Protection Regulation (1.3.1)
GERAM	Generalized Enterprise Reference Architecture and Methodology (2.7.6)
GLBA	Gramm-Leach-Bliley Act (1.3.1)
GRC	governance, risk and compliance (1.3.5)
HIPAA	Health Insurance Portability and Accountability Act (1.3.1)
HITECH	Health Information Technology for Economic and Clinical Health (2.6.1)
HR	human resources (1.4.12)
HSM	hardware security module (1.4.9)
IaaS	infrastructure as a service (1.3.3)
IAASB	International Auditing and Assurance Standards Board (2.10.2)
IAF	Information Assurance Framework (2.6.2)
IAM	identity and access management (2.7.6)
ICT	information and communications technology (1.4.12)
IEC	International Electrotechnical Commission (1.3.1)
IEEE	Institute of Electrical and Electronics Engineers (2.6.2)
IETF	Internet Engineering Task Force (2.6.2)
IIA	Institute of Internal Auditors (4.5.1)
IMDA	Infocomm Media Development Authority (2.6.2)
InfoSrv	information services (2.7.6)
InfraSrv	infrastructure services (2.7.6)
IPPF	International Professional Practices Framework (5.3.1)
IRP	incident response plan (1.4.13)
ISACA	Information Systems Audit and Controls Association (1.3.1)
ISAE	International Standard on Assurance Engagements (2.10.2)
ISO	International Organization for Standardization (1.3.1)
IT	information technology (1.2)
ITAF	Information Technology Audit Framework (5.3.1)
ITAR	International Traffic in Arms Regulations (2.8.3)
ITIL	Information Technology Infrastructure Library (3.3.4)
ITOS	Information technology Operation Services (2.7.6)
ITU-T	International Telecommunications Union–Telecommunication Standardization Sector (2.6.2)
J-SOX	Japan–Sarbanes-Oxley Act, also known as Financial Instruments and Exchange Act (2.6.1)
KPI	key performance indicator (1.4.8)
MFA	multifactor authentication (2.8.3)
NERC	North American Electric Reliability Corporation (3.3.4)
NERC CIP	North American Electric Reliability Corporation Critical Infrastructure Protection (3.3.4)
NIS Directive	Network and Information Systems Directive (2.6.1)
NIST	National Institute of Standards and Technology (USA) (1.4.8)
NIST CSF	National Institute of Standards and Technology (USA) Cybersecurity Framework (1.4.10)
NZISM	New Zealand Information Security Manual (3.7.1)
OASIS	Organization for the Advancement of Structured Information Standards (2.6.2)
OAuth2	Industry-standard protocol for authorization (6.6.2)
OCF	Open Certification Framework (2.6.2)
OCTAVE	Operationally Critical Threat, Asset and Vulnerability Evaluation (1.4.10)
ODCA	Open Data Center Alliance (3.7.1)
OGF	Open Grid Forum (2.6.2)
OS	operating system (1.4.8)
OWASP	Open Web Application Security Project (1.4.14)

PaaS	platform as a service (1.3.3)
PCAOB	Public Company Accounting Oversight Board (1.4.15)
PCI SSC	Payment Card Industry Security Standards Council (1.4.8)
PCI-DSS	Payment Card Industry Data Security Standard (2.6.2)
PHI	protected health information (1.4.10)
PII	personally identifiable information (1.4.12)
PIPEDA	Personal Information Protection and Electronic Documents Act (2.5.2)
PKI	public key infrastructure (1.4.1)
PLA	privacy level agreement (2.6.2)
PR	pull request (8.3.5)
Prod	Production (8.5.1)
PS	presentation services (2.7.6)
QA	quality assurance (8.3)
QOS	quality of service (1.4.1)
RACI	responsible accountable, consulted and informed (1.3.5)
RAM	random access memory (1.4.10)
RBAC	role-based access control (1.4.12)
RTM	requirements traceability matrix (2.7.6)
SaaS	software as a service (1.3.3)
SABSA	Sherwood Applied Business Security Architecture (2.7.6)
SAST	static application security testing (8.4.1)
SCA	software composition analysis (8.4.1)
SCC-CtoP	Standard Contractual Clauses Controller to Processor (1.4.13)
SDLC	software development life cycle (3.6.1)
SDO	standard development organizations (2.6.2)
SDP	software defined perimeter (2.6.2)
SecaaS	security as a service (2.4.6)
SIEM	security information and event management (1.4.9)
SLA	service level agreement (1.3.3)
SLO	service level objectives (1.4.13)
SME	subject matter expert (6.6.4)
SOC	Service Organization Control (1.3.8)
SOX	Sarbanes-Oxley Act (3.4.3)
SP	Special Publication (2.6.2)
SQL	Structured Query Language (1.4.14)
SQO	service qualitative objectives (1.4.13)
SRM	Security and Risk Management (2.7.6)
SSAE	Statement on Standards for Attestation Engagements (2.10.2)
SSH	Secure Shell (2.7.3)
STAR	Security Trust Assurance and Risk (1.3.4)
TOGAF	The Open Group Application Framework (2.7.6)
ToS	terms of service (1.4.13)
TSC	Trust Services Criteria (2.6.2)
UK	United Kingdom (1.3.9)
UI	user interface (5.8.1)
US CLOUD	United States Clarifying Lawful Overseas Use of Data (2.6.1)
UX	user experience (5.8.1)

Page intentionally left blank

Glossary

A

Acceptable use policy—Set of rules applied by the owner, creator or administrator of a network, website, or service, that restrict the ways in which the network, website or system may be used and sets guidelines as to how it should be used

Source: Wikipedia, "Acceptable use policy," en.wikipedia.org/wiki/Acceptable_use_policy

Accountability—The ability to map a given activity or event back to the responsible party

Administrative controls—Refer to policies, procedures or guidelines that define personnel or business practices in accordance with the organization's security goals

Source: Walkowski, D.; "What Are Security Controls?," F5, Inc., 22 August 2019, www.f5.com/labs/articles/education/what-are-security-controls

Agile—The ability to create and respond to change. It is a way of dealing with, and ultimately succeeding in, an uncertain and turbulent environment.

Source: Agile Alliance "Agile 101," www.agilealliance.org/agile101/

AICPA TSC 2017—Trust Services Criteria for security, availability, processing integrity, fonfidentiality and privacy

Source: Association of International Certified Professional Accountants, "TSP Section 100: 2017 Trust Services Criteria for Security, Availability, Processing Integrity, Confidentiality, and Privacy, Includes March 2020 updates," www.aicpa.org/content/dam/aicpa/interestareas/frc/assuranceadvisoryservices/downloadabledocuments/trust-services-criteria-2020.pdf

Algorithm—A mathematical function that is used in the encryption and decryption processes

Source: (ICS)^2, "CISSP Glossary - Student Guide," www.isc2.org/Certifications/CISSP/CISSP-Student-Glossary#

Ansible—An open-source software provisioning, configuration management, and application-deployment tool enabling infrastructure as code

Source: Wikipedia, "Ansible (software)," en.wikipedia.org/wiki/Ansible_(software)

Application layer—In the Open Systems Interconnection (OSI) communications model, the application layer provides services for an application program to ensure that effective communication with another application program in a network is possible.

Assessments—The process of identifying risks to organizational operations (including mission, functions, image, reputation), organizational assets, individuals, other organizations, and the Nation, resulting from the operation of a system—Generally, the purpose of an assessment is to get a snapshot of the current reality of your organization

Source: Lean Learning Center, "Audit & Assessment Services," https://leanlearningcenter.com/services/audit-and-assessment/

Assurance—Pursuant to an accountable relationship between two or more parties, an IT audit and assurance professional is engaged to issue a written communication expressing a conclusion about the subject matters for which the accountable party is responsible. Assurance refers to a number of related activities designed to provide the reader or user of the report with a level of assurance or comfort over the subject matter.

Attack Vector—A path or route used by the adversary to gain access to the target (asset)

Attestation—CPA issues a report in which he or she expresses an opinion or a conclusion on the subject matter (for example, financial statements) so that a user can make informed decisions.

Source: Association of International Certified Professional Accountants, "CPA Services," 2015, www.aicpa.org/Advocacy/State/DownloadableDocuments/Attest-Services-Chart-Color-Nonfillable.pdf

Attestations—Legal acknowledgment of the authenticity of a document and a verification that proper processes were followed (e.g., SOC2 report from CSP, 3rd Party independent assessment on CSP)

Source: Halton, C.; "Attestation," Investopedia, 17 July 2019, www.investopedia.com/terms/a/attestation.asp

Audit—1) An audit is focused on compliance. An audit also measures the current reality, but it then compares it against a specific standard.

Source: Lean Learning Center, "Audit & Assessment Services," https://leanlearningcenter.com/services/audit-and-assessment/

2) Systematic, independent and documented process for obtaining audit evidence and evaluating it objectively to determine the extent to which the audit criteria are fulfilled

Source: ISO, ISO 9000:2015 *Quality management systems – Fundamentals and vocabulary*, September 2015, www.iso.org/standard/45481.html

Audit Scope—Extent and boundaries of an audit

Source: ISO, ISO 19011:2018 *Guidelines for auditing management systems*

Auditing—The independent assessment conducted by a qualified assessor of the conformity of the internal and external (cloud) processes within the scope of the applicable regulatory requirements, organizational policies and/or standard requirements

Availability—Property of being accessible and usable upon demand by an authorized entity

Source: ISO, ISO/IEC 17788:2014(en) *Information technology – Cloud computing – Overview and vocabulary*, www.iso.org/obp/ui/#iso:std:iso-iec:17788:ed-1:v1:en

B

Botnet—A botnet is a number of Internet-connected devices, each of which is running one or more bots. Botnets can be used to perform Distributed Denial-of-Service (DDoS) attacks, steal data, send spam, and allows the attacker to access the device and its connection.

Source: Wikipedia, "Botnet," en.wikipedia.org/wiki/Botnet

Breach—The loss of control, compromise, unauthorized disclosure, unauthorized acquisition, or any similar occurrence where: a person other than an authorized user accesses or potentially accesses personally identifiable information

Business continuity (BC)—Preventing, mitigating and recovering from disruption: The terms 'business resumption planning', 'disaster recovery planning' and 'contingency planning' also may be used in this context; they focus on recovery aspects of continuity, and for that reason the 'resilience' aspect should also be taken into account.

Business Continuity Planning (BCP)—Business continuity planning (BCP) involves planning and procedural aspects, encompassing emergency response, crisis communications, business continuity and disaster recovery.

Source: Lingeswara Satyanarayana Tammineed, R.; "Key Issues, Challenges and Resolutions in Implementing Business Continuity Projects," *ISACA Journal*,1 January 2012, www.isaca.org/resources/isaca-journal/past-issues/2012/key-issues-challenges-and-resolutions-in-implementing-business-continuity-projects

C

C5 BSI—A baseline security level for cloud computing. It is used by professional cloud service providers, auditors and cloud customers.

Source: National Institute of Standards and Technology, "Security and Privacy Controls for Information Systems and Organizations," NIST Special Publication 800-53, Revision 5, September 2020, nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r5.pdf

Capabilities—Reinforcing security and privacy controls implemented by technical, physical, and procedural means

Source: Bundesamt für Sicherheit in der Informationstechnik, "Criteria Catalogue C5," www.bsi.bund.de/EN/Topics/CloudComputing/Compliance_Criteria_Catalogue/Compliance_Criteria_Catalogue_node.html

Certification—The provision by an independent body of written assurance (a certificate) that the product, service or system in question meets specific requirements

Source: ISO. "Certification," https://www.iso.org/certification.html

Chain of custody—A legal principle regarding the validity and integrity of evidence. It requires accountability for anything that will be used as evidence in a legal proceeding to ensure that it can be accounted for from the time it was collected until the time it is presented in a court of law.

Scope Notes: Includes documentation as to who had access to the evidence and when, as well as the ability to identify evidence as being the exact item that was recovered or tested. Lack of control over evidence can lead to it being discredited. Chain of custody depends on the ability to verify that evidence could not have been tampered with. This is accomplished by sealing off the evidence, so it cannot be changed, and providing a documentary record of custody to prove that the evidence was at all times under strict control and not subject to tampering.

Chef—A configuration management tool

Source: Wikipedia, "Chef (software)," en.wikipedia.org/wiki/Chef_(software)

CI/CD Pipeline—A series of steps that involves continuous automation and monitoring to deliver new versions of software. The steps that form a CI/CD pipeline are distinct subsets of tasks that typically include build, test, release, deploy, and validate.

Source: Red Hat, Inc., "What is a CI/CD pipeline?," www.redhat.com/en/topics/devops/what-cicd-pipeline

Ciphertext—Information generated by an encryption algorithm to protect the plaintext and that is unintelligible to the unauthorized reader

Source: (ICS)², "CISSP Glossary - Student Guide," www.isc2.org/Certifications/CISSP/CISSP-Student-Glossary#

Cloud access security broker—On-premises or cloud-based software that sits between cloud service users and cloud applications, and monitors all activity and enforces security policies

Source: Wikipedia, "Cloud access security broker," en.wikipedia.org/wiki/Cloud_access_security_broker

Cloud auditor—Cloud service partner with the responsibility to conduct an audit of the provision and use of cloud services

Source: ISO, ISO/IEC 17788:2014(en) *Information technology – Cloud computing – Overview and vocabulary*, 2014, www.iso.org/obp/ui/#iso:std:iso-iec:17788:ed-1:v1:en

Cloud computing—Paradigm for enabling network access to a scalable and elastic pool of shareable physical or virtual resources with self-service provisioning and administration on-demand

Source: ISO, ISO/IEC 17788:2014(en) *Information technology – Cloud computing – Overview and vocabulary*, www.iso.org/obp/ui/#iso:std:iso-iec:17788:ed-1:v1:en

Cloud customer—A person or organization that is a customer of a cloud; note that a cloud customer may itself be a cloud and that clouds may offer services to one another

Source: National Institute of Standards and Technology, "Cloud Computing Synopsis and Recommendations," Special Publication 800-146, May 2012, nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-146.pdf

Cloud database—A database accessible to clients from the cloud and delivered to users on demand via the Internet. Also referred to as Database as a Service (DBaaS), cloud databases can use cloud computing to achieve optimized scaling, high availability, multi-tenancy, and effective resource allocation

Source: Gordon, A.; *The Official (ISC)2 Guide to the CCSP CBK*, 2nd Edition, Amazon, www.amazon.com/Official-ISC-Guide-CCSP-CBK/dp/1119276721

Cloud service customer—Party which is in a business relationship for the purpose of using cloud services

Source: ISO, ISO/IEC 17788:2014(en) *Information technology – Cloud computing – Overview and vocabulary*, 2014, www.iso.org/obp/ui/#iso:std:iso-iec:17788:ed-1:v1:en

Cloud service provider—Party which makes cloud services available

Source: ISO, ISO/IEC 17788:2014(en) *Information technology – Cloud computing – Overview and vocabulary*, 2014, www.iso.org/obp/ui/#iso:std:iso-iec:17788:ed-1:v1:en

Cloud-native computing—Cloud-native computing is a development methodology that assumes the entire context of cloud technology; applications are designed for, and anticipated to scale within, public, private and hybrid clouds. Characteristic technologies include containers, microservices, serverless functions and immutable infrastructure.

Source: Wikipedia, "Cloud Native Computing," https://en.wikipedia.org/wiki/Cloud_native_computing

COBIT—1. COBIT 2019: The current iteration of COBIT builds on and integrates more than 25 years of developments in the field of enterprise governance of information and technology (I&T), not only incorporating new insights from science, but also operationalizing these insights as practices. COBIT is a broad and comprehensive I&T governance and management framework and continues to establish itself as a generally accepted framework for I&T governance.

Scope Notes: Earlier versions of COBIT focused on IT, whereas COBIT 2019 focuses on information **and** technology aimed at the whole enterprise, recognizing that I&T has become crucial in the support, sustainability and growth of enterprises. (See www.isaca.org/cobit for more information.)

2. COBIT 5: Formerly known as Control Objectives for Information and related Technology (COBIT); with this iteration used only as the acronym. A complete, internationally accepted framework for governing and managing enterprise information and technology (IT) that supports enterprise executives and management in their definition and achievement of business goals and related IT goals. COBIT describes five principles and seven enablers that support enterprises in the development, implementation and continuous improvement and monitoring of good IT-related governance and management practices.

Scope Notes: Earlier versions of COBIT focused on control objectives related to IT processes, management and control of IT processes and IT governance aspects. Adoption and use of the COBIT framework are supported by guidance from a growing family of supporting products.

Community cloud—The cloud infrastructure is provisioned for exclusive use by a specific community of consumers from organizations that have shared concerns (e.g., mission, security requirements, policy, and compliance considerations). It may be owned, managed, and operated by one or more of the organizations in the community, a third party, or some combination of them, and it may exist on or off premises.

Source: National Institute of Standards and Technology, "The NIST Definition of Cloud Computing," Special Publication 800-145, September 2011, nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf

Compensating control—An internal control that reduces the risk of an existing or potential control weakness resulting in errors and omissions

Compliance—Adherence to, and the ability to demonstrate adherence to, mandated requirements defined by laws and regulations, as well as voluntary requirements resulting from contractual obligations and internal policies

Source: ISO, ISO/IEC 17788:2014(en) *Information technology – Cloud computing – Overview and vocabulary*, 2014, www.iso.org/obp/ui/#iso:std:iso-iec:17788:ed-1:v1:en

Confidentiality—Property that information is not made available or disclosed to unauthorized individuals, entities, or processes

Source: ISO, ISO/IEC 17788:2014(en) *Information technology – Cloud computing – Overview and vocabulary*, 2014, www.iso.org/obp/ui/#iso:std:iso-iec:17788:ed-1:v1:en

Configuration management—A systems engineering process for establishing and maintaining consistency of a product's performance, functional, and physical attributes with its requirements, design, and operational information throughout its life

Source: Wikipedia, "Configuration management," en.wikipedia.org/wiki/Configuration_management

Container—A method for packaging and securely running an application within an application virtualization environment. Also known as an application container or a server application container

Source: National Institute of Standards and Technology, "Application Container Security Guide," NIST Special Publication 800-190, September 2017, nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-190.pdf

Continuous assurance/compliance—The combination of continuous auditing and continuous monitoring

Source: Cipriano , H.M.; R. Pereira; R. Almeida; M. Mira da Silva; *Addressing Continuous Auditing Challenges in the Digital Age: A Literature Review*, 2019, www.igi-global.com/chapter/addressing-continuous-auditing-challenges-in-the-digital-age/222624

Continuous audit—An on-going assessment process that aims to determine the fulfilment of Service Qualitative Objectives (SQOs) and Service Level Objectives (SLOs), conducted at a frequency requested by the purpose of audit

Source: Cloud Security Alliance, "European Security Certification Framework," Version 1.0, Project Number 731845, EUSEC D1.4 Principles, Criteria and Requirements for a Multi-Party Recognition and Continuous Auditing Based Certifications, March 2018, cdn0.scrvt.com/fokus/c56a737828fef8cf/79add0dec99a/D1.4-Multiparty-recognition-V1.0.pdf

Continuous delivery—The extension of Continuous Integration that involves the automated testing and release of validated code to a repository resulting in a codebase that is always ready for deployment to a production environment

Source: Red Hat, Inc., "What is CI/CD?," www.redhat.com/en/topics/devops/what-is-ci-cd

Continuous deployment—The final phase of an automated software release process that uses scripts or tools to move code changes in a repository to production after it passes automated testing

Source: Red Hat, Inc., "What is CI/CD?," www.redhat.com/en/topics/devops/what-is-ci-cd

Continuous integration—The first phase of an automated software release process that involves using automated tools to frequently build, test, and merge code changes from multiple contributors into a single software project

Source: Atlassian, "What is Continuous Integration?," www.atlassian.com/continuous-delivery/continuous-integration

Continuous monitoring—Maintaining ongoing awareness of information security, vulnerabilities, and threats to support organization risk management decision

Source: National Institute of Standards and Technology, "Information Security Continuous Monitoring (ISCM) for Federal Information Systems and Organizations," NIST Special Publication 800-137, September 2011, nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-137.pdf

Contracts—A binding legal agreement between two or more parties that may be enforceable in a court of law

Source: National Institute of Standards and Technology, "Protecting Controlled Unclassified Information in Nonfederal Systems and Organizations," NIST.SP.800-171 Rev. 2, February 2020, csrc.nist.gov/publications/detail/sp/800-171/rev-2/final

Control framework—A set of fundamental controls that facilitates the discharge of business process owner responsibilities to prevent financial or information loss in an enterprise

Control objective—A statement of the desired result or purpose to be achieved by implementing control procedures in a particular process

Controls—Controls are intended to reduce the frequency or impact of realized risk.

Corporate governance—The system by which enterprises are directed and controlled. The board of directors is responsible for the governance of their enterprise. It consists of the leadership and organizational structures and processes that ensure the enterprise sustains and extends strategies and objectives.

Corrective—Include any measures taken to repair damage or restore resources and capabilities to their prior state following an unauthorized or unwanted activity

Source*:* Walkowski, D.; "What Are Security Controls?," F5, Inc., 22 August 2019, www.f5.com/labs/articles/education/what-are-security-controls

Credit risk—A credit risk is risk of default on a debt that may arise from a borrower failing to make required payments.

Source: Wikipedia, "Credit risk," en.wikipedia.org/wiki/Credit_risk

Cryptographic algorithm—A cryptographic checksum is created by performing a complicated series of mathematical operations (known as a cryptographic algorithm) that translates the data in the file into a fixed string of digits called a hash value, which is then used as the checksum.

Source: (ICS)^2, "CISSP Glossary – Student Guide," www.isc2.org/Certifications/CISSP/CISSP-Student-Glossary#

CSA Security Guidance—The fourth version of the Security Guidance for Critical Areas of Focus in Cloud Computing is built on previous iterations of the security guidance, dedicated research, and public participation from the Cloud Security Alliance members, working groups, and the industry experts within our community. This version incorporates advances in cloud, security, and supporting technologies; reflects on real-world cloud security practices; integrates the latest Cloud Security Alliance research projects; and offers guidance for related technologies.

Source: Cloud Security Alliance, "CSA Security Guidance for Critical Areas of Focus in Cloud Computing v4.0," https://cloudsecurityalliance.org/research/guidance/

D

Detective—Describe any security measure taken or solution that's implemented to detect and alert to unwanted or unauthorized activity in progress or after it has occurred

Source: Walkowski, D.; "What Are Security Controls?," F5, Inc., 22 August 2019, www.f5.com/labs/articles/education/what-are-security-controls

DevOps—DevOps is a set of practices that combines software development (Dev) and IT operations (Ops). It aims to shorten the system development life cycle and provide continuous delivery with high software quality. DevOps is complementary with Agile software development; several DevOps aspects came from Agile methodology.

Source: Wikipedia, "DevOps," en.wikipedia.org/wiki/DevOps

DevSecOps—An augmentation of DevOps to allow for the integration of security practices in the DevOps approach

Source: Wikipedia, "DevOps," https://en.wikipedia.org/wiki/DevOps

Digital signature—A piece of information, a digitized form of signature, that provides sender authenticity, message integrity and non-repudiation. A digital signature is generated using the sender's private key or applying a one-way hash function.

Disaster Recovery (DR)—Disaster recovery (DR) is the technical component of BCP and focuses on the continuity of information and communication technology systems that support business functions.

Source: Tammineed, R.L.S.; "Key Issues, Challenges and Resolutions in Implementing Business Continuity Projects," *ISACA Journal*, 1 January 2012, www.isaca.org/resources/isaca-journal/past-issues/2012/key-issues-challenges-and-resolutions-in-implementing-business-continuity-projects

Distributed denial-of-service attack (DDoS)—A denial-of-service (DoS) assault from multiple sources

Due care—The level of care expected from a reasonable person of similar competency under similar conditions

Source: (ICS)², "CISSP Glossary - Student Guide," www.isc2.org/Certifications/CISSP/CISSP-Student-Glossary#

Due diligence—The performance of those actions that are generally regarded as prudent, responsible and necessary to conduct a thorough and objective investigation, review and/or analysis

Dynamic application security testing—A set of tools used to test software during operation and provide feedback on compliance and general security issues. DAST tools are typically used during the testing and QA phase.

Source: Imperva, "SAST, DAST, IAST and RASP," www.imperva.com/learn/application-security/sast-iast-dast/

E

Encryption—The process of transforming plaintext into ciphertext using a cryptographic algorithm and key

Source: National Institute of Standards and Technology, "Recommendation for Pair-Wise Key-Establishment Using Integer Factorization Cryptography," NIST SP 800-56B Rev. 2, March 2019, https://csrc.nist.gov/publications/detail/sp/800-56b/rev-2/final

Enterprise—An organization with a defined mission/goal and a defined boundary, using systems to execute that mission, and with responsibility for managing its own risks and performance

Source: National Institute of Standards and Technology, "Security and Privacy Controls for Information Systems and Organizations," NIST Special Publication 800-53 Rev. 5, September 2020, https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r5.pdf

Enterprise architecture—Description of the fundamental underlying design of the components of the business system, or of one element of the business system (e.g., technology), the relationships among them, and the manner in which they support the enterprise's objectives

Enterprise Architecture Model—A methodology and a set of tools that enable security architects, enterprise architects and risk management professionals to leverage a common set of solutions that fulfill their common needs to be able to assess where their internal IT and their cloud providers are in terms of security capabilities and to plan a roadmap to meet the security needs of their business

Source: Cloud Security Alliance, "EAWG™ Enterprise Architecture," https://ea.cloudsecurityalliance.org/

Enterprise risk management—A process, effected by an entity's board of directors, management and other personnel, applied in strategy setting and across the enterprise, designed to identify potential events that may affect the entity, and manage risk to be within its risk appetite, to provide reasonable assurance

Source: Committee of Sponsoring Organizations of the Treadway Commission, "Enterprise Risk Management —Integrated Framework" Executive Summary," September 2004, www.coso.org/Documents/COSO-ERM-Executive-Summary.pdf

Event—Any observable occurrence in a network or information system

Source: National Institute of Standards and Technology, "Risk Management Framework for Information Systems and Organizations: A System Life Cycle Approach for Security and Privacy," SP 800-37 Rev. 2, December 2018, https://csrc.nist.gov/publications/detail/sp/800-37/rev-2/final

F

Face value—The value of a coin, stamp or paper money as printed on the coin, stamp or bill itself by the issuing authority

Source: Wikipedia, "Face value," https://en.wikipedia.org/wiki/Face_value

Financial risk—Is any of various types of risk associated with financing, including financial transactions that include company loans in risk of default

Source: Wikipedia, "Financial risk," https://en.wikipedia.org/wiki/Financial_risk

FIPS 140-3—This standard specifies the security requirements that will be satisfied by a cryptographic module utilized within a security system protecting sensitive but unclassified information.

Source: National Institute of Standards and Technology, "Security Requirements for Cryptographic Modules," Information Technology Laboratory, 22 March 2019, https://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.140-3.pdf

Food and Drug Administration (FDA)—The Food and Drug Administration (FDA) is responsible for protecting the public health by assuring the safety, efficacy, and security of human and veterinary drugs, biological products, medical devices, our nation's food supply, cosmetics, and products that emit radiation.

Source: www.usa.gov/federal-agencies/food-and-drug-administration

Framework—Provides a common organizing structure for multiple approaches by assembling standards, guidelines, and practices that are working effectively today.

Source: National Institute of Standards and Technology, "Approaches for Federal Agencies to Use the Cybersecurity Framework," NISTIR 8170, March 2020, *https://csrc.nist.gov/publications/detail/nistir/8170/final*

G

General Data Protection Regulation (GDPR)—The General Data Protection Regulation (GDPR) is a regulation in EU law on data protection and privacy in the European Union (EU) and the European Economic Area (EEA). It also addresses the transfer of personal data outside the EU and EEA areas.

Source: Wikipedia, "General Data Protection Regulation," https://en.wikipedia.org/wiki/General_Data_Protection_Regulation

Governance—Ensures that stakeholder needs, conditions and options are evaluated to determine balanced, agreed-on enterprise objectives to be achieved; setting direction through prioritization and decision making; and monitoring performance and compliance against agreed-on direction and objectives

Governance framework—A framework is a basic conceptual structure used to solve or address complex issues. An enabler of governance. A set of concepts, assumptions and practices that define how something can be approached or understood, the relationships amongst the entities involved, the roles of those involved, and the boundaries (what is and is not included in the governance system).

Gramm–Leach–Bliley Act (GLBA)—The Gramm–Leach–Bliley Act (GLBA), also known as the Financial Services Modernization Act of 1999, removing barriers in the market among banking companies, securities companies and insurance companies that prohibited any one institution from acting as any combination of an investment bank, a commercial bank, and an insurance company

Source: Wikipedia, "*Gramm–Leach–Bliley_Act,"* https://en.wikipedia.org/wiki/Gramm–Leach–Bliley_Act

H

Health Insurance Portability and Accountability Act of 1996 (HIPAA)—The Health Insurance Portability and Accountability Act of 1996 (HIPAA or the Kennedy–Kassebaum Act was enacted by the 104th United States Congress and signed by President Bill Clinton in 1996. It was created primarily to modernize the flow of healthcare information, stipulate how personally identifiable information maintained by the healthcare and healthcare insurance industries should be protected from fraud and theft, and address limitations on healthcare insurance coverage.

Source: Wikipedia, "Health Insurance Portability and Accountability Act," https://en.wikipedia.org/wiki/Health_Insurance_Portability_and_Accountability_Act

Hybrid cloud—1) The cloud infrastructure is a composition of two or more distinct cloud infrastructures (private, community, or public) that remain unique entities, but are bound together by standardized or proprietary technology that enables data and application portability (e.g., cloud bursting for load balancing between clouds).

Source: National Institute of Standards and Technology, "The NIST Definition of Cloud Computing," Special Publication 800-145, September 2011, https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf

2) A type of cloud computing that combines on-premises infrastructure (private cloud) with a public cloud

Source: Microsoft Azure, "What are public, private, and hybrid clouds?," https://azure.microsoft.com/en-us/overview/what-are-private-public-hybrid-clouds/#overview

I

Identity and access management (IAM)—A framework of policies and technologies for ensuring that the proper people in an enterprise have the appropriate access to technology resources

Source: Wikipedia, "Identity management," https://en.wikipedia.org/wiki/Identity_management

Incident—An occurrence that results in actual or potential jeopardy to the confidentiality, integrity, or availability of an information system or the information the system processes, stores, or transmits or that constitutes a violation or imminent threat of violation of security policies, security procedures, or acceptable use policies

Source: National Institute of Standards and Technology, "Computer Security Incident Handling Guide," Special Publication 800-61 Rev. 2, August 2012, https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-61r2.pdf

Incident response—The mitigation of violations of security policies and recommended practices

Source: National Institute of Standards and Technology, Computer Security Resource Center, *Glossary*, https://csrc.nist.gov/glossary/term/incident_response

Incident response plan—The documentation of a predetermined set of instructions or procedures to detect, respond to, and limit consequences of a malicious cyber attacks against an organization's information systems(s)

Source: National Institute of Standards and Technology, Computer Security Resource Center, *Glossary*, https://csrc.nist.gov/glossary/term/incident_response_plan

Information security—Preservation of confidentiality, integrity, and availability of information

Source: ISO, ISO/IEC 17788:2014(en) *Information technology – Cloud computing – Overview and vocabulary*, 2014, https://www.iso.org/obp/ui/#iso:std:iso-iec:17788:ed-1:v1:en

Infrastructure as a Service (IaaS)—The capability provided to the consumer is to provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications

Source*:* National Institute of Standards and Technology, "The NIST Definition of Cloud Computing," Special Publication 800-145, September 2011, https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf

Inherent risk—The risk level or exposure without taking into account the actions that management has taken or might take (e.g., implementing controls)

Initialization vector (IV)—A non-secret binary vector used as the initializing input algorithm, or a random starting point, for the encryption of a plaintext block sequence to increase security by introducing additional cryptographic variance and to synchronize cryptographic equipment

Source: (ICS)^2, "CISSP Glossary - Student Guide," www.isc2.org/Certifications/CISSP/CISSP-Student-Glossary#

Integrity—Property of accuracy and completeness

Source: ISO, ISO/IEC 17788:2014(en) *Information technology – Cloud computing – Overview and vocabulary*, 2014, www.iso.org/obp/ui/#iso:std:iso-iec:17788:ed-1:v1:en

Internet of Things (IoT)—Network of physical objects—"things"—that are embedded with sensors, software, and other technologies for the purpose of connecting and exchanging data with other devices and systems over the Internet

Source: Wikipedia, "Internet of things," https://en.wikipedia.org/wiki/Internet_of_things

Interoperability—A characteristic of a product or system, whose interfaces are completely understood, to work with other products or systems, at present or in the future, in either implementation or access, without any restrictions

Source: Wikipedia, "Interoperability," https://en.wikipedia.org/wiki/Interoperability

ISO/IEC 22301—Requirements to plan, establish, implement, operate, monitor, review, maintain and continually improve a documented management system to protect against, reduce the likelihood of occurrence, prepare for, respond to, and recover from disruptive incidents when they arise

Source: ISO, ISO 22301:2019 *Security and resilience – Business continuity management systems – Requirements*, 2019, *www.iso.org/standard/75106.html*

ISO/IEC 27001—Requirements for an information security management system

Source: ISO, ISO/IEC 27001 *Information Security Management*, www.iso.org/isoiec-27001-information-security.html

ISO/IEC 27002—Guidelines for organizational information security standards and information security management practices including the selection, implementation and management of controls

Source: ISO, ISO/IEC 27002:2013 *Information technology – Security techniques – Code of practice for information security controls*, 2013, www.iso.org/standard/54533.html

ISO/IEC 27017—Code of practice for information security controls

Source: ISO, ISO/IEC 27017:2013 *Information technology – Security techniques – Code of practice for information security controls*, 2013, www.iso.org/standard/43757.html

ISO/IEC 27018—Code of practice for protection of personally identifiable information (PII) in public clouds acting as PII processors

Source: ISO, ISO/IEC 27018:2019 *Information technology – Security techniques – Code of practice for protection of personally identifiable information (PII) in public clouds acting as PII processors*, 2019, www.iso.org/standard/76559.html

ISO/IEC 27035—Information security incident management

Source: ISO, ISO/IEC 27035-1:2016 *Information technology – Security techniques – Information security incident management – Part 1: Principles of incident management*, 2016, www.iso.org/standard/60803.html

ISO/IEC 38500—Information technology – Governance of IT for the organization

Source: ISO, ISO/IEC 38500:2015 *Information technology – Governance of IT for the organization*, 2015, www.iso.org/standard/62816.html

IT governance framework—A model that integrates a set of guidelines, policies and methods that represent the organizational approach to IT governance

J

Jericho Forum—An international IT security thought-leadership group dedicated to defining ways to deliver effective IT security solutions

Source: Cloud Security Alliance, "Jericho Forum and Cloud Security Alliance join forces to address Cloud Computing Security," 27 May 2009, https://cloudsecurityalliance.org/press-releases/2009/05/27/jericho-forum-and-cloud-security-alliance-join-forces-to-address-cloud-computing-security/

Judgmental sampling—Any sample that is selected subjectively or in such a manner that the sample selection process is not random or the sampling results are not evaluated mathematically.

Jurisdictions—Authority granted to a legal body to administer justice, as defined by the kind of case, and the location of the issue

Source: Wikipedia, "Jurisdiction," https://en.wikipedia.org/wiki/Jurisdiction

K

Key management—Dealing with the generation, exchange, storage, use, crypto-shredding (destruction) and replacement of keys

Source: Wikipedia, "Key management," https://en.wikipedia.org/wiki/Key_management

L

Legacy environment—Environments that are on premises of the organization and not in the cloud

Legal risk—Legal risk is the risk of loss to an institution which is primarily caused by: (a) a defective transaction; or (b) a claim (including a defense to a claim or a counterclaim) being made or some other event occurring which results in a liability for the institution or other loss (for example, as a result of the termination of a contract) or; (c) failing to take appropriate measures to protect assets (for example, intellectual property) owned by the institution; or (d) change in law

Source: Wikipedia, "Legal risk," https://en.wikipedia.org/wiki/Legal_risk

M

Malware—Software or firmware intended to perform an unauthorized process that will have adverse impact on the confidentiality, integrity, or availability of an information system. A virus, worm, Trojan horse, or other code-based entity that infects a host. Spyware and some forms of adware are also examples of malicious code.

Source: National Institute of Standards and Technology, "Security and Privacy Controls for Information Systems and Organizations," NIST Special Publication 800-53, Revision 5, September 2020, https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r5.pdf

Maturity—Indicates the degree of reliability or dependency or capability that the business can place on a process achieving the desired goals or objectives

Source: ISACA, "Appraisals," https://cmmiinstitute.com/learning/appraisals/levels

N

NIST SP800-137—Information Security Continuous Monitoring (ISCM) for Federal Information Systems and Organizations

Source: National Institute of Standards and Technology, "Information Security Continuous Monitoring (ISCM) for Federal Information Systems and Organizations," NIST Special Publication 800-137, September 2011, https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-137.pdf

NIST SP800-53—Security and Privacy Controls for Federal Information Systems and Organizations

Source: National Institute of Standards and Technology, "Security and Privacy Controls for Information Systems and Organizations," NIST Special Publication 800-53, Revision 5, September 2020, https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r5.pdf

Non-statistical sampling—Method of selecting a portion of a population, by means of own judgement and experience, for the purpose of quickly confirming a proposition. This method does not allow drawing mathematical conclusions on the entire population.

O

On premises—On-premises software (commonly misstated as on-premise, and alternatively abbreviated "on-prem") is installed and runs on computers on the premises of the person or organization using the software, rather than at a remote facility such as a server farm or cloud.

Source: Wikipedia, "On-premises software," https://en.wikipedia.org/wiki/On-premises_software

P

Patch management—An area of systems management that involves acquiring, testing and installing multiple patches (code changes) to an administered computer system in order to maintain up-to-date software and often to address security risk

PCAOB AS5—Public Company Accounting Oversight Board is adopting Auditing Standard No. 5, An Audit of Internal Control Over Financial Reporting That Is Integrated with An Audit of Financial Statements, as well as an independence rule and conforming amendments to the Board's auditing standards.

Source: Public Company Accounting Oversight Board, "Auditing Standard No. 5 An Audit of Internal Control Over Financial Reporting That Is Integrated with An Audit of Financial Statements," https://pcaobus.org/oversight/standards/archived-standards/pre-reorganized-auditing-standards-interpretations/details/Auditing_Standard_5

PCI DSS—The Payment Card Industry Data Security Standard (PCI DSS) is an information security standard for organizations that handle branded credit cards from the major card schemes.

Source: Wikipedia, "Payment Card Industry Data Security Standard," https://en.wikipedia.org/wiki/Payment_Card_Industry_Data_Security_Standard

Penetration testing—A method of testing where testers target individual binary components or the application as a whole to determine whether intra or intercomponent vulnerabilities can be exploited to compromise the application, its data, or its environment resources

Source: National Institute of Standards and Technology, Computer Security Resource Center, *Glossary*, https://csrc.nist.gov/glossary/term/penetration_testing

Phishing—A technique for attempting to acquire sensitive data, such as bank account numbers, or access to a larger computerized system through a fraudulent solicitation in email or on a web site. The perpetrator typically masquerades as a legitimate business or reputable person.

Source: National Institute of Standards and Technology, Small Business Cybersecurity Corner, "Glossary," www.nist.gov/itl/smallbusinesscyber/cybersecurity-basics/glossary

Physical controls—Describe anything tangible that's used to prevent or detect unauthorized access to physical areas, systems, or assets

Source: Walkowski, D.; "What Are Security Controls?," F5, Inc., 22 August 2019, www.f5.com/labs/articles/education/what-are-security-controls

Plaintext—The message in its natural format has not been turned into a secret.

Source: (ICS)², "CISSP Glossary - Student Guide," www.isc2.org/Certifications/CISSP/CISSP-Student-Glossary#

Platform as a Service (PaaS)—The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the provider.

Source: National Institute of Standards and Technology, "The NIST Definition of Cloud Computing," Special Publication 800-145, September 2011, https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf

Policies—Capture the intent, direction and expectation of Management

Source: ISACA, *Information Security Governance: Guidance for Information Security Managers*, 2008, www.isaca.org/bookstore/it-governance-and-business-management/w3itg

Policy—Generally, a document that records a high-level principle or course of action that has been decided on

Portability—The ability of a computer program to be ported from one system to another

Source: Wikipedia, "Portability," https://en.wikipedia.org/wiki/Portability

Preventative—Describe any security measure that's designed to stop unwanted or unauthorized activity from occurring.

Source: Walkowski, D.; "What Are Security Controls?," F5, Inc., 22 August 2019, www.f5.com/labs/articles/education/what-are-security-controls

Private cloud—1) The cloud infrastructure is provisioned for exclusive use by a single organization comprising multiple consumers (e.g., business units). It may be owned, managed, and operated by the organization, a third party, or some combination of them, and it may exist on or off premises.

Source: National Institute of Standards and Technology, "The NIST Definition of Cloud Computing," Special Publication 800-145, September 2011, https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf

2) A cloud deployment model in which the services and infrastructure are always maintained on a private network, and the hardware and software are dedicated solely to one organization. A private cloud can be hosted in an organization's on-site data center or by a third-party service provider.

Source: Microsoft Azure, "What are public, private, and hybrid clouds?," https://azure.microsoft.com/en-us/overview/what-are-private-public-hybrid-clouds/#overview

Procedure—A document containing a detailed description of the steps necessary to perform specific operations in conformance with applicable standards. Procedures are defined as part of processes.

Procedures—Process that includes necessary steps to accomplish a specific goal

ISACA, *Information Security Governance: Guidance for Information Security Managers*, 2008, www.isaca.org/bookstore/it-governance-and-business-management/w3itg

Process—Set of interrelated or interacting activities which transforms inputs into outputs

Source: ISO 9000 *Introduction and Support Package: Guidance on the Concept and Use of the Process Approach for management systems*, ISO/TC 176/SC 2/N 544R3, www.iso.org/files/live/sites/isoorg/files/archive/pdf/en/04_concept_and_use_of_the_process_approach_for_management_systems.pdf

Provider lock-in / Vendor lock-in—A customer dependent on a vendor for products and services, unable to use another vendor without substantial switching costs

Source: Wikipedia, "Vendor lock-in," https://en.wikipedia.org/wiki/Vendor_lock-in

Proxy—An application that "breaks" the connection between client and server

Source: National Institute of Standards and Technology, Computer Security Resource Center, *Glossary*, https://csrc.nist.gov/glossary/term/proxy

Proxy server—A server that acts on behalf of a user

Source: (ICS)², "CISSP Glossary - Student Guide," www.isc2.org/Certifications/CISSP/CISSP-Student-Glossary#

Public cloud—1) The cloud infrastructure is provisioned for open use by the general public. It may be owned, managed, and operated by a business, academic, or government organization, or some combination of them. It exists on the premises of the cloud provider.

Source: National Institute of Standards and Technology, "The NIST Definition of Cloud Computing," Special Publication 800-145, September 2011, https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf

2) A cloud deployment model in which all hardware, software, and other supporting infrastructure are owned and managed by a third-party cloud provider and are shared with other organizations, or cloud tenants

Source: Microsoft Azure, "What are public, private, and hybrid clouds?," https://azure.microsoft.com/en-us/overview/what-are-private-public-hybrid-clouds/#overview

Puppet—Configuration management software

Source: Wikipedia, "Puppet (company)," https://en.wikipedia.org/wiki/Puppet_(company)

R

RACI-style matrix—Illustrates who is Responsible, Accountable, Consulted and Informed within an organizational framework

Source: (ICS)², "CISSP Glossary - Student Guide," www.isc2.org/Certifications/CISSP/CISSP-Student-Glossary#

Ransomware —A type of malware that attempts to deny access to a user's data, usually by encrypting the data with a key known only to the hacker who deployed the malware, until a ransom is paid

Source: National Institute of Standards and Technology, Small Business Cybersecurity Corner, "Glossary," www.nist.gov/itl/smallbusinesscyber/cybersecurity-basics/glossary

Regression test—Re-running functional and non-functional tests to ensure that previously developed and tested software still performs after a change

Source: Wikipedia, "Regression testing," https://en.wikipedia.org/wiki/Regression_testing

Regulation—Rules or laws defined and enforced by an authority to regulate conduct

Regulatory risk—Regulatory risk is the risk that a change in laws and regulations will materially impact a security, business, sector, or market.

Source: Investopedia, "Regulatory Risk," www.investopedia.com/terms/r/regulatory_risk.asp

Remediation—After vulnerabilities are identified and assessed, appropriate remediation can take place to mitigate or eliminate the vulnerability.

Reputation risk—The current and prospective effect on earnings and capital arising from negative public opinion. Reputation risk affects a bank's ability to establish new relationships or services, or to continue servicing existing relationships. It may expose the bank to litigation, financial loss or a decline in its customer base. A bank's reputation can be damaged by Internet banking services that are executed poorly or otherwise alienate customers and the public. An Internet bank has a greater reputation risk as compared to a traditional brick-and-mortar bank, because it is easier for its customers to leave and go to a different Internet bank and since it cannot discuss any problems in person with the customer.

Residual risk—1) The remaining risk after management has implemented a risk response

2) The ability of an information system to operate under adverse conditions or stress, even if in a degraded or debilitated state, while maintaining essential operational capabilities, and to recover to an effective operational posture in a time frame consistent with mission needs

Source: National Institute of Standards and Technology, "Security and Privacy Controls for Information Systems and Organizations," NIST Special Publication 800-53, Revision 5, September 2020, https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r5.pdf

Resilience—1) The ability of an information system to operate under adverse conditions or stress, even if in a degraded or debilitated state, while maintaining essential operational capabilities, and to recover to an effective operational posture in a time frame consistent with mission needs

Source: National Institute of Standards and Technology, "Security and Privacy Controls for Information Systems and Organizations," NIST Special Publication 800-53, Revision 5, September 2020, https://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r5.pdf

2) The ability to prepare for and adapt to changing conditions and withstand and recover rapidly from disruptions. Resilience includes the ability to withstand and recover from deliberate attacks, accidents, or naturally occurring threats or incidents.

Source: National Institute of Standards and Technology, Computer Security Resource Center, *Glossary*, https://csrc.nist.gov/glossary/term/resilience

Risk—Effect of uncertainty on objectives

Source: ISO, ISO 31000:2018(en) *Risk management – Guidelines*, 2018, www.iso.org/obp/ui/#iso:std:iso:31000:ed-2:v1:en

Risk appetite—The amount of risk, on a broad level, that an entity is willing to accept in pursuit of its mission

Risk assessment—A process used to identify and evaluate risk and its potential effects

Risk management—The coordinated activities to direct and control an enterprise with regard to risk

Risk profile—The amount of risk that is involved in an investment.

Source: Cambridge Dictionary, "risk profile," https://dictionary.cambridge.org/es-LA/dictionary/english/risk-profile

Risk register—A repository of the key attributes of potential and known IT risk issues. Attributes may include name, description, owner, expected/actual frequency, potential/actual magnitude, potential/actual business impact, disposition.

Risk tolerance—The acceptable level of variation that management is willing to allow for any particular risk as the enterprise pursues its objectives

S

Salt Stack—Python-based, open-source software for event-driven IT automation, remote task execution, and configuration management

Source: Wikipedia, "Salt (software)," https://en.wikipedia.org/wiki/Salt_(software)

Sandbox—Is a testing environment that isolates untested code changes and outright experimentation from the production environment or repository, in the context of software development including Web development and revision control.

Source: Wikipedia, "Sandbox," https://en.wikipedia.org/wiki/Sandbox_(software_development)

Sarbanes-Oxley (SOX)—"Corporate and Auditing Accountability, Responsibility, and Transparency Act" (in the House) and more commonly called Sarbanes-Oxley, Sarbox or SOX, is a United States federal law that set new or expanded requirements for all U.S. public company boards, management and public accounting firms. A number of provisions of the Act also apply to privately held companies, such as the willful destruction of evidence to impede a federal investigation.

Source: Wikipedia, "Sarbanes-Oxley Act,"

Security as a Service (SecaaS)—The next generation of managed security services dedicated to the delivery, over the Internet, of specialized information-security services

Security program—The overall combination of technical, operational and procedural measures and management structures implemented to provide for the confidentiality, integrity and availability of information based on business requirements and risk analysis

Security risk—Something or someone likely to cause danger or difficulty

Source: Cambridge Dictionary, "security risk," https://dictionary.cambridge.org/dictionary/english/security-risk

Serverless computing—A flexible "pay-as-you-go" cloud computing execution model in which the cloud provider runs the server and dynamically manages the allocation of machine resources. Pricing is based on the amount of actual resources consumed by an application, so the developers pay only for the backend services they use.

Source: Wikipedia, "Serverless computing," https://en.wikipedia.org/wiki/Serverless_computing

Service level agreement—Documented agreement between the service provider and customer that identifies services and service targets

Source: ISO, ISO/IEC 17788:2014(en) *Information technology – Cloud computing – Overview and vocabulary*, 2014, www.iso.org/obp/ui/#iso:std:iso-iec:17788:ed-1:v1:en

Service level agreement (SLA)—An agreement, preferably documented, between a service provider and the customer(s)/user(s) that defines minimum performance targets for a service and how they will be measured

Shadow IT—Refers to IT devices, software and services outside the ownership or control of IT organizations

Source: Gartner, *Gartner Glossary*, "Shadow IT," www.gartner.com/en/information-technology/glossary/shadow

Shared responsibility model—The compliance responsibility between the cloud customer and the cloud service provider based on the degree of control each party has over the architecture stack

Social engineering—An attack based on deceiving users or administrators at the target site into revealing confidential or sensitive information

Software as a Service (SaaS)—The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface.

Source: National Institute of Standards and Technology, "The NIST Definition of Cloud Computing," Special Publication 800-145, September 2011, https://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf

Software development life cycle (SDLC)—Software development life cycle (SDLC) is composed of the phases deployed in the development or acquisition of a software system.

Spearphishing—A highly targeted phishing attack, usually to a specific individual or department within an organization

Source: National Institute of Standards and Technology, Small Business Cybersecurity Corner, "Glossary," www.nist.gov/itl/smallbusinesscyber/cybersecurity-basics/glossary

Stakeholders—Anyone who has a responsibility for, an expectation from or some other interest in the enterprise

Standards—Metrics, allowable boundaries or the process used to determine whether procedures meet policy requirements

Source: ISACA, *Information Security Governance: Guidance for Information Security Managers*, 2008, www.isaca.org/bookstore/it-governance-and-business-management/w3itg

STAR Program—The Security Trust Assurance and Risk (STAR) Program encompasses key principles of transparency, rigorous auditing, and harmonization of standards.

Source: Cloud Security Alliance, "CSA Security Trust Assurance and Risk (STAR)," https://cloudsecurityalliance.org/star/

Static application security testing—A set of technologies designed to analyze application source code, byte code and binaries for coding and design conditions that are indicative of security vulnerabilities. SAST solutions are typically used during the development phase.

Source: Gartner, *Gartner Glossary*, Static Application Security Testing (SAST), www.gartner.com/en/information-technology/glossary/static-application-security-testing-sast

Statistical sampling—A method of selecting a portion of a population, by means of mathematical calculations and probabilities, for the purpose of making scientifically and mathematically sound inferences regarding the characteristics of the entire population

Status quo—The existing state of affairs

Source: Merriam-Webster, "status quo," www.merriam-webster.com/dictionary/status%20quo#synonyms

Subsidiary—A company wholly controlled by another

Source: Merriam-Webster, "subsidiary," www.merriam-webster.com/dictionary/subsidiary

T

Technical controls—(Also known as logical controls) include hardware or software mechanisms used to protect assets

Source: Walkowski, D.; "What Are Security Controls?," F5, Inc., 22 August 2019, www.f5.com/labs/articles/education/what-are-security-controls

Testware—A sub-set of software with a special purpose, that is, for software testing, especially for software testing automation

Source: Wikipedia, "Testware," https://en.wikipedia.org/wiki/Testware

Third party—1) An outside source from the internal company

2) A third person or organization less directly involved in a matter than the main people or organizations that are involved

Source: Cambridge Dictionary, "third party," https://dictionary.cambridge.org/dictionary/english/third-party

Threat—Any circumstance or event with the potential to adversely impact organizational operations (including mission, functions, image, or reputation), organizational assets, or individuals through an information system via unauthorized access, destruction, disclosure, modification of information, and/or denial of service

Source: National Institute of Standards and Technology, Computer Security Resource Center, *Glossary*, https://csrc.nist.gov/glossary/term/threat

Threat modeling—A process by which potential threats or the absence of appropriate safeguards, can be identified, enumerated, and mitigations can be prioritized

Source: Wikipedia, "Threat model," https://en.wikipedia.org/wiki/Threat_model

Turtle diagram—A process auditing and improvement tool (schematic visual representation) of the key elements that make up a process including the resources needed to achieve the process's purpose.

V

Virtualization—The simulation of the software and/or hardware upon which other software runs

Source: National Institute of Standards and Technology, Computer Security Resource Center, *Glossary,* https://csrc.nist.gov/glossary/term/Virtualization

Vulnerability management—An Information Security Continuous Monitoring (ISCM) capability that identifies vulnerabilities [Common Vulnerabilities and Exposures (CVEs)] on devices that are likely to be used by attackers to compromise a device and use it as a platform from which to extend compromise to the network

Source: National Institute of Standards and Technology, Computer Security Resource Center, *Glossary,* https://csrc.nist.gov/glossary/term/Vulnerability_Management

Vulnerability —A weakness in the design, implementation, operation or internal control of a process that could expose the system to adverse threats from threat events

W

Web application—Application software that runs on a web server, unlike computer-based software programs that are stored locally on the Operating System (OS) of the device

Source: Wikipedia, "Web application," https://en.wikipedia.org/wiki/Web_application