DOMAIN 1

Architectural Concepts and Design Requirements

THE GOAL OF THE Architectural Concepts and Design Requirements domain is to provide you with knowledge of the building blocks necessary to develop cloud-based systems.

You will be introduced to such cloud computing concepts as the customer, provider, partner, measured services, scalability, virtualization, storage, and networking. You will be able to understand the cloud reference architecture based on activities defined by industry-standard documents.

Lastly, you will gain knowledge in relevant security and design principles for cloud computing, including secure data lifecycle and cost-benefit analysis of cloud-based systems.

DOMAIN OBJECTIVES

After completing this domain, you will be able to do the following:

- ❑ **Define the various roles, characteristics, and technologies as they relate to cloud computing concepts**
- ❑ **Describe cloud computing concepts as they relate to cloud computing activities, capabilities, categories, models, and cross-cutting aspects**
- ❑ **Identify the design principles necessary for secure cloud computing**
- ❑ **Define the various design principles for the different types of cloud categories**
- ❑ **Describe the design principles for secure cloud computing**
- ❑ **Identify criteria specific to national, international, and industry for certifying trusted cloud services**
- ❑ **Identify criteria specific to the system and subsystem product certification**

INTRODUCTION

> *"Cloud computing is a model for enabling ubiquitous, convenient, on-demand network access to a shared pool of configurable computing resources (e.g., networks, servers, storage, applications, and services) that can be rapidly provisioned and released with minimal management effort or service provider interaction."*
>
> "The NIST Definition of Cloud Computing"[1]

Cloud computing (*Figure 1.1*) is the use of Internet-based computing resources, typically "as a service," to allow internal or external customers to consume where scalable and elastic information technology (IT)-enabled capabilities are provided.

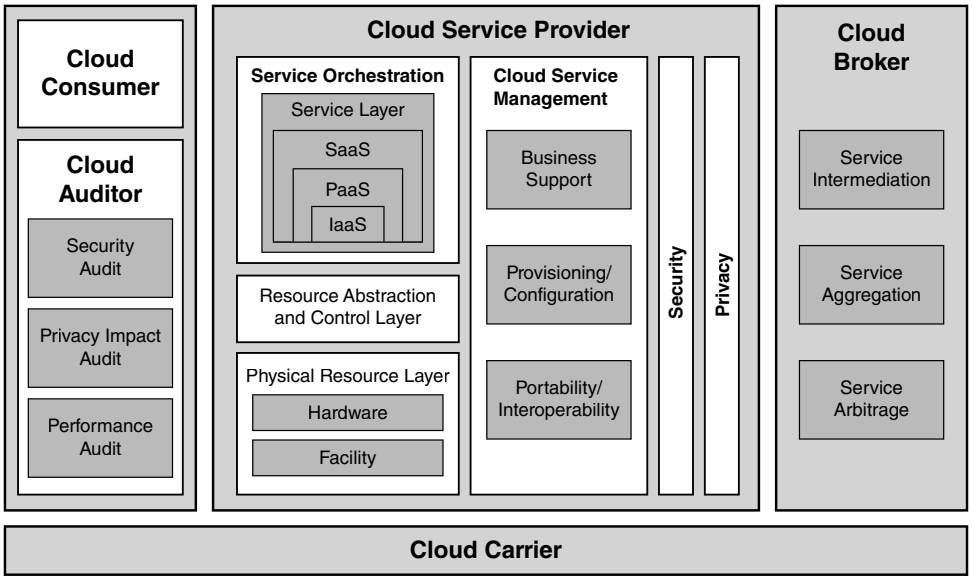

FIGURE 1.1 **Cloud computing overview.**

Cloud computing, or cloud, means many things to many people. There are indeed various definitions for cloud computing and what it means from many of the leading standards bodies. The previous National Institute of Standards and Technology (NIST) definition is the most commonly utilized, cited by professionals and others alike to clarify what the term *cloud* means.

It's important to note the difference between a cloud service provider (CSP) and a managed service provider (MSP). The main difference is to be found in the control exerted over the data and process and by who. With an MSP, the consumer dictates the

technology and operating procedures. According to the MSP Alliance, MSPs typically have the following distinguishing characteristics:[2]

- Some form of network operations center (NOC) service
- Some form of help desk service
- Remote monitoring and management of all or most of the objects for the customer
- Proactive maintenance of the objects under management for the customer
- Delivery of these solutions with some form of predictable billing model, where the customer knows with great accuracy what the regular IT management expense will be

With a CSP, the service provider dictates both the technology and the operational procedures being made available to the cloud consumer. This means that the CSP is offering some or all of the components of cloud computing through a software as a service (SaaS), infrastructure as a service (IaaS), or platform as a service (PaaS) model.

Drivers for Cloud Computing

There are many drivers that may move a company to consider cloud computing. These may include the costs associated with the ownership of their current IT infrastructure solutions as well as projected costs to continue to maintain these solutions year in and year out (*Figure 1.2*).

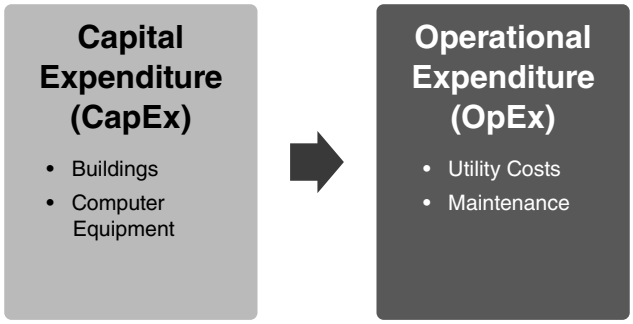

FIGURE 1.2 Drivers that move companies toward cloud computing.

Additional drivers include but are not limited to the following:

- The desire to reduce IT complexity
  - **Risk reduction:** Users can use the cloud to test ideas and concepts before making major investments in technology.
  - **Scalability:** Users have access to a large number of resources that scale based on user demand.

- **Elasticity:** The environment transparently manages a user's resource utilization based on dynamically changing needs.

- Consumption-based pricing
  - **Virtualization:** Each user has a single view of the available resources, independent of their arrangement in terms of physical devices.
  - **Cost:** The pay-per-usage model allows an organization to pay only for the resources it needs with basically no investment in the physical resources available in the cloud. There are no infrastructure maintenance or upgrade costs.
- Business agility
  - **Mobility:** Users can access data and applications from around the globe.
  - **Collaboration and innovation:** Users are starting to see the cloud as a way to work simultaneously on common data and information.

Security, Risks, and Benefits

You cannot bring up or discuss the topic of cloud computing without hearing the words *security*, *risk*, and *compliance*. In truth, cloud computing does pose challenges and represents a paradigm shift in the way in which technology solutions are being delivered. As with any notable change, this brings about questions and a requirement for clear and concise understandings and interpretations to be obtained, from both a customer and a provider perspective. The Certified Cloud Security Professional (CCSP) must play a key role in the dialogue within the organization as it pertains to cloud computing, its role, the opportunity costs, and the associated risks (*Figure 1.3*).

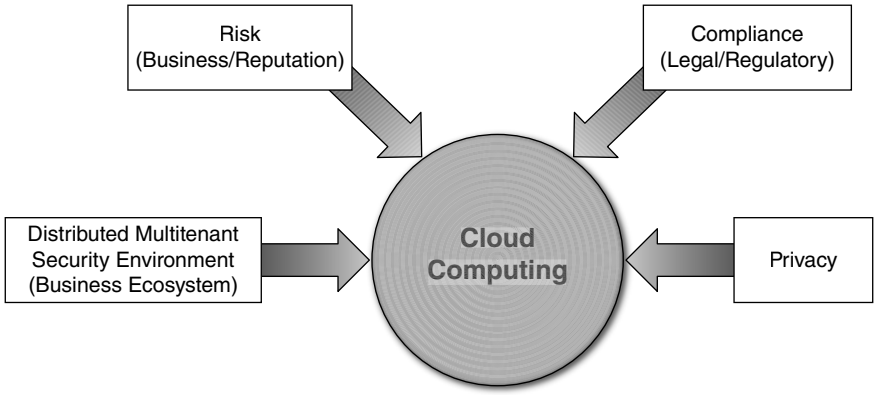

FIGURE 1.3 **Cloud computing issues and concerns.**

Risk can take many forms in an organization. The organization needs to carefully weigh all the risks associated with a business decision before engaging in an activity to

minimize the risk impact associated with an activity. There are many approaches and frameworks that can be used to address risk in an organization, such as the Control Objectives for Information and Related Technology (COBIT) framework, The Committee of Sponsoring Organizations of the Treadway Commission (COSO) Enterprise Risk Management Integrated Framework, and the NIST Risk Management Framework. Organizations need to become risk aware in general, focusing on risks within and around the organization that may cause harm to the reputation of the business. Reputational risk can be defined as "the loss of value of a brand or the ability of an organization to persuade."[3] To manage reputational risk, an organization should consider the following items:

- Strategic alignment
  - Effective board oversight
  - Integration of risk into strategy setting and business planning
- Cultural alignment
  - Strong corporate values and a focus on compliance
- Operational focus
  - Strong control environment

Although many people think of cloud technologies as less secure or carrying greater risk, this is simply not possible or acceptable to say unless making a direct and measured comparison against a specified environment or service. For instance, it would be incorrect to simply assume or state that cloud computing is less secure as a service modality for the delivery of a customer relationship management (CRM) platform than a more traditional CRM application model, calling for an on-premise installation of the CRM application and its supporting infrastructure and databases. To assess the true level of security and risk associated with each model of ownership and consumption, the two platforms would need to be compared across a range of factors and issues, allowing for a side-by-side comparison of the key deliverables and issues associated with each model.

In truth, the cloud may be more or less secure than your organization's environment and current security controls depending on any number of factors, which include technological components; risk management processes; preventative, detective, and corrective controls; governance and oversight processes; resilience and continuity capabilities; defense in depth; and multifactor authentication.

Therefore, the approach to security varies depending on the provider and the ability for your organization to alter and amend its overall security posture prior to, during, and after migration or utilization of cloud services.

In the same way that no two organizations or entities are the same, neither are two CSPs. A one-size-fits-all approach is never good for security, so do not settle for it when utilizing cloud-based services.

The extensive use of automation within the cloud enables real-time monitoring and reporting on security control points, allowing for the establishment of continuous security monitoring regimes, enhancing the overall security posture of the organization consuming the cloud services. The benefits realized by the organization can include greater security visibility, enhanced policy and governance enforcement, and a better framework for management of the extended business ecosystem through a transition from an infrastructure-centric to a data-centric security model.

CLOUD COMPUTING DEFINITIONS

The following list forms a common set of terms and phrases you will need to become familiar with as a CCSP. Having an understanding of these items puts you in a strong position to communicate and understand technologies, deployments, solutions, and architectures within the organization as needed. This list is not comprehensive and should be used along with the vocabulary terms in Appendix B, "Glossary," to form as complete a picture as possible of the language of cloud computing.

- **Anything as a service (XaaS):** The growing diversity of services available over the Internet via cloud computing as opposed to being provided locally or on premises.
- **Apache CloudStack:** An open source cloud computing and IaaS platform developed to help make creating, deploying, and managing cloud services easier by providing a complete stack of features and components for cloud environments.
- **Business continuity:** The capability of the organization to continue delivery of products or services at acceptable predefined levels following a loss of service.
- **Business continuity management:** A holistic management process that identifies potential threats to an organization and the impacts to business operations those threats, if realized, might cause. It provides a framework for building organizational resilience with the capability of an effective response that safeguards the interests of its key stakeholders, reputation, brand, and value-creating activities.
- **Business continuity plan:** The creation of a strategy through the recognition of threats and risks facing a company, with an eye to ensure that personnel and assets are protected and able to function in the event of a disaster.
- **Cloud app:** Short for cloud application, cloud app describes a software application that is never installed on a local computer. Instead, it is accessed via the Internet.
- **Cloud Application Management for Platforms (CAMP):** CAMP is a specification designed to ease management of applications—including packaging and deployment—across public and private cloud computing platforms.

- **Cloud backup:** Cloud backup, or cloud computer backup, refers to backing up data to a remote, cloud-based server. As a form of cloud storage, cloud backup data is stored in and accessible from multiple distributed and connected resources that comprise a cloud.
- **Cloud backup solutions:** Cloud backup solutions enable enterprises or individuals to store their data and computer files on the Internet using a storage service provider rather than storing the data locally on a physical disk, such as a hard drive or tape backup.
- **Cloud computing:** A type of computing, comparable to grid computing, that relies on sharing computing resources and using a network of remote servers to store, manage, and process data instead of using a local server or a personal computer.
- **Cloud computing accounting software:** Cloud computing accounting software is accounting software that is hosted on remote servers. It provides accounting capabilities to businesses in a fashion similar to the SaaS business model. Data is sent into the cloud, where it is processed and returned to the user. All application functions are performed offsite, not on the user's desktop.
- **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, multitenancy, and effective resource allocation.
- **Cloud enablement:** The process of making available one or more of the following services and infrastructures to create a public cloud computing environment: CSP, client, and application.
- **Cloud management:** Software and technologies designed for operating and monitoring the applications, data, and services residing in the cloud. Cloud management tools help ensure a company's cloud computing–based resources are working optimally and properly interacting with users and other services.
- **Cloud migration:** The process of transitioning all or part of a company's data, applications, and services from onsite premises behind the firewall to the cloud, where the information can be provided over the Internet on an on-demand basis.
- **Cloud OS:** A phrase frequently used in place of PaaS to denote an association to cloud computing.
- **Cloud portability:** In cloud computing terminology, this refers to the ability to move applications and their associated data between one CSP and another—or between public and private cloud environments.

- **Cloud provisioning:** The deployment of a company's cloud computing strategy, which typically first involves selecting which applications and services will reside in the public cloud and which will remain onsite behind the firewall or in the private cloud. Cloud provisioning also entails developing the processes for interfacing with the cloud's applications and services as well as auditing and monitoring who accesses and utilizes the resources.
- **Cloud server hosting:** A type of hosting in which hosting services are made available to customers on demand via the Internet. Rather than being provided by a single server or virtual server, cloud server hosting services are provided by multiple connected servers that comprise a cloud.
- **Cloud storage:** The storage of data online in the cloud, whereby a company's data is stored in and accessible from multiple distributed and connected resources that comprise a cloud.
- **Cloud testing:** Load and performance testing conducted on the applications and services provided via cloud computing—particularly the capability to access these services—to ensure optimal performance and scalability under a variety of conditions.
- **Desktop as a service:** A form of virtual desktop infrastructure (VDI) in which the VDI is outsourced and handled by a third party. Also called hosted desktop services, desktop as a service is frequently delivered as a cloud service along with the apps needed for use on the virtual desktop.
- **Enterprise application:** Describes applications—or software—that a business uses to assist the organization in solving enterprise problems. When the word *enterprise* is combined with *application*, it usually refers to a software platform that is too large and complex for individual or small business use.
- **Enterprise cloud backup:** Enterprise-grade cloud backup solutions typically add essential features such as archiving and disaster recovery (DR) to cloud backup solutions.
- **Eucalyptus:** An open source cloud computing and IaaS platform for enabling AWS-compatible private and hybrid clouds.
- **Event:** A change of state that has significance for the management of an IT service or other configuration item. The term can also be used to mean an alert or notification created by an IT service, configuration item, or monitoring tool. Events often require IT operations staff to take actions and lead to incidents being logged.
- **Host:** A device providing a service.
- **Hybrid cloud storage:** A combination of public cloud storage and private cloud storage in which some critical data resides in the enterprise's private cloud and other data is stored and accessible from a public cloud storage provider.

- **IaaS:** IaaS is defined as computer infrastructure, such as virtualization, being delivered as a service. IaaS is popular in the data center where software and servers are purchased as a fully outsourced service and usually billed on usage and how much of the resource is used—compared with the traditional method of buying software and servers outright.
- **Incident:** An unplanned interruption to an IT service or reduction in the quality of an IT service.
- **Managed service provider:** An IT service provider in which the customer dictates both the technology and the operational procedures.
- **Mean time between failure (MTBF):** The measure of the average time between failures of a specific component or part of a system.
- **Mean time to repair (MTTR):** The measure of the average time it should take to repair a failed component or part of a system.
- **Mobile cloud storage**: A form of cloud storage that applies to storing an individual's mobile device data in the cloud and providing the individual with access to the data from anywhere.
- **Multitenant:** In cloud computing, multitenant is the phrase used to describe multiple customers using the same public cloud.
- **Node:** A physical connection.
- **Online backup:** In storage technology, online backup means to back up data from your hard drive to a remote server or computer using a network connection. Online backup technology leverages the Internet and cloud computing to create an attractive offsite storage solution with few hardware requirements for any business of any size.
- **PaaS:** The process of deploying onto the cloud infrastructure consumer-created or acquired applications that are created using programming languages, libraries, services, and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems (OSs), or storage but has control over the deployed applications and possibly the configuration settings for the application-hosting environment.
- **Personal cloud storage:** A form of cloud storage that applies to storing an individual's data in the cloud and providing the individual with access to the data from anywhere. Personal cloud storage also often enables syncing and sharing stored data across multiple devices such as mobile phones and tablet computers.
- **Private cloud:** Describes a cloud computing platform that is implemented within the corporate firewall, under the control of the IT department. A private cloud

is designed to offer the same features and benefits of cloud systems but removes a number of objections to the cloud computing model, including control over enterprise and customer data, worries about security, and issues connected to regulatory compliance.

- **Private cloud project:** Companies initiate private cloud projects to enable their IT infrastructure to become more capable of quickly adapting to continually evolving business needs and requirements. Private cloud projects can also be connected to public clouds to create hybrid clouds.
- **Private cloud security:** A private cloud implementation aims to avoid many of the objections regarding cloud computing security. Because a private cloud setup is implemented safely within the corporate firewall, it remains under the control of the IT department.
- **Private cloud storage:** A form of cloud storage in which both the enterprise data and the cloud storage resources reside within the enterprise's data center and behind the firewall.
- **Problem:** The unknown cause of one or more incidents, often identified as a result of multiple similar incidents.
- **Public cloud storage:** A form of cloud storage in which the enterprise and storage service provider are separate and the data is stored outside of the enterprise's data center.
- **Recovery point objective (RPO):** The RPO helps determine how much information must be recovered and restored. Another way of looking at the RPO is to ask yourself, "How much data can the company afford to lose?"
- **Recovery time objective (RTO):** A time measure of how fast you need each system to be up and running in the event of a disaster or critical failure.
- **SaaS:** A software delivery method that provides access to software and its functions remotely as a web-based service. SaaS allows organizations to access business functionality at a cost typically less than paying for licensed applications since SaaS pricing is based on a monthly fee.
- **Storage cloud:** Refers to the collection of multiple distributed and connected resources responsible for storing and managing data online in the cloud.
- **Vertical cloud computing:** Describes the optimization of cloud computing and cloud services for a particular vertical (for example, a specific industry) or specific-use application.
- **Virtual host:** A software implementation of a physical host.

CLOUD COMPUTING ROLES

The following groups form the key roles and functions associated with cloud computing. They do not constitute an exhaustive list but highlight the main roles and functions within cloud computing:

- **Cloud backup service provider:** A third-party entity that manages and holds operational responsibilities for cloud-based data backup services and solutions to customers from a central data center.
- **Cloud computing reseller:** A company that purchases hosting services from a cloud server hosting or cloud computing provider and then resells them to its own customers.
- **Cloud customer:** An individual or entity that utilizes or subscribes to cloud-based services or resources.
- **Cloud service auditor**: A third-party organization that verifies attainment of service-level agreements (SLAs).
- **Cloud services brokerage (CSB):** Typically a third-party entity or company that looks to extend or enhance value to multiple customers of cloud-based services through relationships with multiple CSPs. It acts as a liaison between cloud services customers and CSPs, selecting the best provider for each customer and monitoring the services. The CSB can be utilized as a middleman to broker the best deal and customize services to the customer's requirements. The CSB may also resell cloud services.
- **CSP:** A company that provides cloud-based platform, infrastructure, application, or storage services to other organizations or individuals, usually for a fee; otherwise known to clients "as a service."

KEY CLOUD COMPUTING CHARACTERISTICS

Think of the following as a rulebook or a set of laws when dealing with cloud computing. *If a service or solution does not meet all of the following key characteristics, it is not true cloud computing.*

- **On-demand self-service:** The cloud service provided that enables the provision of cloud resources on demand (whenever and wherever they are required). From a security perspective, this has introduced challenges to governing the use and provisioning of cloud-based services, which may violate organizational policies.

By its nature, on-demand self-service does not require procurement, provisioning, or approval from finance, and as such, it can be provisioned by almost anyone with a credit card. For enterprise customers, this is most likely the least important characteristic because self-service for the majority of end users is not of utmost importance.

- **Broad network access:** The cloud, by its nature, is an always on and always accessible offering for users to have widespread access to resources, data, and other assets. Think convenience—access what you want, when you need it, from any location.

  In theory, all you should require is Internet access and relevant credentials and tokens, which give you access to the resources.

  The mobile device and smart device revolution that is altering the way organizations fundamentally operate has introduced an interesting dynamic into the cloud conversation within many organizations. These devices should also be able to access the relevant resources that a user may require; however, compatibility issues, the inability to apply security controls effectively, and nonstandardization of platforms and software systems has stemmed this somewhat.

- **Resource pooling:** Lies at the heart of all that is good about cloud computing. More often than not, traditional, noncloud systems may see utilization rates for their resources between 80 percent and 90 percent for a few hours a week and rates at an average of 10 percent to 20 percent for the remainder. What the cloud looks to do is group (pool) resources for use across the user landscape or multiple clients, which can then scale and adjust to the user's or client's needs, based on their workload or resource requirements. CSPs typically have large numbers of resources available, from hundreds to thousands of servers, network devices, applications, and so on, which can accommodate large volumes of customers and can prioritize and facilitate appropriate resourcing for each client.

- **Rapid elasticity:** Allows the user to obtain additional resources, storage, compute power, and so on, as the user's need or workload requires. This is more often transparent to the user, with more resources added as necessary seamlessly.

  Because cloud services utilize the pay-per-use concept, you pay for what you use. This is of particular benefit to seasonal or event-type businesses utilizing cloud services.

  Think of a provider selling 100,000 tickets for a major sporting event or concert. Leading up to the ticket release date, little to no compute resources are needed; however, when the tickets go on sale, they may need to accommodate 100,000 users in the space of 30–40 minutes. This is where rapid elasticity and cloud

computing can really be beneficial, compared with traditional IT deployments, which would have to invest heavily using capital expenditure (CapEx) to support such demand.

- **Measured service:** Cloud computing offers a unique and important component that traditional IT deployments have struggled to provide—resource usage can be measured, controlled, reported, and alerted upon, which results in multiple benefits and overall transparency between the provider and the client. In the same way you may have a metered electricity service or a mobile phone that you top up with credit, these services allow you to control and be aware of costs. Essentially, you pay for what you use and have the ability to get an itemized bill or breakdown of usage.

  A key benefit being availed by many proactive organizations is the ability to charge departments or business units for their use of services, thus allowing IT and finance to quantify exact usage and costs per department or by business function—something that was incredibly difficult to achieve in traditional IT environments.

In theory and in practice, cloud computing should have large resource pools to enable swift scaling, rapid movement, and flexibility to meet your needs at any given time within the bounds of your service subscription.

Without all these characteristics, it is simply not possible for the user to be confident and assured that the delivery and continuity of services will be maintained in line with potential growth or sudden scaling (either upward or downward). Without pooling and measured services, you cannot implement the cloud computing economic model.

CLOUD TRANSITION SCENARIO

Consider the following scenario.

Due to competitive pressures, XYZ Corp is hoping to better leverage the economic and scalable nature of cloud computing. These policies have driven XYZ Corp toward the consideration of a hybrid cloud model that consists of enterprise private and public cloud use. Although security risk has driven many of the conversations, a risk management approach has allowed the company to separate its data assets into two segments: sensitive and nonsensitive. IT governance guidelines must now be applied across the entire cloud platform and infrastructure security environment. This also affects infrastructure operational options. XYZ Corp must now apply cloud architectural concepts and design requirements that would best align with corporate business and security goals.

As a CCSP, you have several issues to address to guide XYZ Corp through its planned transition to a cloud architecture.

- What cloud deployment model(s) would need to be assessed to select the appropriate ones for the enterprise architecture?
  - **a.** Based on the choice(s) made, additional issues may become apparent, such as these:
    - **i.** Who will the audiences be?
    - **ii.** What types of data will they be using and storing?
    - **iii.** How will secure access to the cloud be enabled, audited, managed, and removed?
    - **iv.** When and where will access be granted to the cloud? Under what constraints (time, location, platform, and so on)?
- What cloud service model(s) would need to be chosen for the enterprise architecture?
  - **a.** Based on the choice(s) made, additional issues may become apparent, such as these:
    - **i.** Who will the audiences be?
    - **ii.** What types of data will they be using and storing?
    - **iii.** How will secure access to the cloud service be enabled, audited, managed, and removed?
    - **iv.** When and where will access be granted to the cloud service? Under what constraints (time, location, platform, and so on)?

Dealing with a scenario such as this requires the CCSP to work with the stakeholders in XYZ Corp to seek answers to the questions posed. In addition, the CCSP should carefully consider the information in Table 1.1 to craft a solution.

TABLE 1.1 Possible Solutions

INFORMATION ITEM	POSSIBLE SOLUTION
Hybrid cloud model	Outsourced hosting in partnership with on-premise IT support
Risk-management-driven data separation	Data classification scheme implemented company wide
IT governance guidelines	Coordination of all governance, risk, and compliance (GRC) activities within XYZ Corp through a chief risk officer (CRO) role
Cloud architecture alignment with business requirements	Requirements gathering and documentation exercise driven by a project management office (PMO) or a business analyst (BA) function

BUILDING BLOCKS

The building blocks of cloud computing are composed of random access memory (RAM), the central processing unit (CPU), storage, and networking. IaaS has the most fundamental building blocks of any cloud service: the processing, storage, and network infrastructure upon which all cloud applications are built. In a typical IaaS scenario, the service provider delivers the server, storage, and networking hardware and its virtualization, and then it's up to the customer to implement the OSs, middleware, and applications required.

CLOUD COMPUTING FUNCTIONS

As with traditional computing and technology environments, a number of functions are essential for creating, designing, implementing, testing, auditing, and maintaining the relevant assets. The same is true for cloud computing, with the following key roles representing a sample of the fundamental components and personnel required to operate cloud environments:

- **Cloud administrator:** This individual is typically responsible for the implementation, monitoring, and maintenance of the cloud within the organization or on behalf of an organization (acting as a third party).

  Most notably, this role involves the implementation of policies, permissions, access to resources, and so on. The cloud administrator works directly with system, network, and cloud storage administrators.

- **Cloud application architect:** This person is typically responsible for adapting, porting, or deploying an application to a target cloud environment.

  The main focus of this role is to work closely and alongside development and other design and implementation resources to ensure that an application's performance, reliability, and security are all maintained throughout the lifecycle of the application. This requires continuous assessment, verification, and testing throughout the various phases of both the software and systems development lifecycles.

  Most architects represent a mix or blend of system administration experience and domain-specific expertise—giving insight to the OS, domain, and other

components, while identifying potential reasons the application may be experiencing performance degradation or other negative impacts.

- **Cloud architect:** This role determines when and how a private cloud meets the policies and needs of an organization's strategic goals and contractual requirements from a technical perspective.

  The cloud architect is also responsible for designing the private cloud, is involved in hybrid cloud deployments and instances, and has a key role in understanding and evaluating technologies, vendors, services, and other skillsets needed to deploy the private cloud or to establish and function the hybrid cloud components.
- **Cloud data architect:** This individual is similar to the cloud architect. The data architect's role is to ensure the various storage types and mechanisms utilized within the cloud environment meet and conform to the relevant SLAs and that the storage components are functioning according to their specified requirements.
- **Cloud developer:** This person focuses on development for the cloud infrastructure itself. This role can vary from client tools or solutions engagements to systems components. Although developers can operate independently or as part of a team, regular interactions with cloud administrators and security practitioners are required for debugging, code reviews, and relevant security assessment remediation requirements.
- **Cloud operator:** This individual is responsible for daily operational tasks and duties that focus on cloud maintenance and monitoring activities.
- **Cloud service manager:** This person is typically responsible for policy design, business agreement, pricing model, and some elements of the SLA (not necessarily the legal components or amendments that require contractual amendments). This role works closely with cloud management and customers to reach agreement and alongside the cloud administrator to implement SLAs and policies on behalf of the customers.
- **Cloud storage administrator:** This role focuses on the mapping, segregations, bandwidth, and reliability of storage volumes assigned. Additionally, this role may require ensuring that conformance to relevant SLAs continues to be met, working with and alongside network and cloud administrators.

CLOUD SERVICE CATEGORIES

Cloud service categories fall into three main groups: IaaS, PaaS, and SaaS. Each is discussed in the following sections.

IaaS

According to "The NIST Definition of Cloud Computing," in 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 OSs and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over OSs, storage, and deployed applications; and possibly limited control of select networking components (e.g., host firewalls)."[4]

Traditionally, infrastructure has always been the focal point for ensuring which capabilities and organization requirements could be met versus those that were restricted. It also represented possibly the most significant investments in terms of CapEx and skilled resources made by the organization. The emergence of the cloud has changed this traditional view of infrastructure's role significantly by commoditizing it and allowing it to be consumed through an on-demand, pay-as-you-go model.

IaaS Key Components and Characteristics

The following form the basis for the IaaS service model:

- **Scale:** The requirement for automation and tools to support the potentially significant workloads of either internal users or those across multiple cloud deployments (dependent on which cloud service offering) is a key component of IaaS. Users and customers require optimal levels of visibility, control, and assurances related to the infrastructure and its ability to satisfy their requirements.
- **Converged network and IT capacity pool:** This follows from the scale focus, but it looks to drill into the virtualization and service management components required to cover and provide appropriate levels of service across network boundaries.

  From a customer or user perspective, the pool appears seamless and endless (no visible barriers or restrictions, along with minimal requirement to initiate additional resources) for both the servers and the network. These are (or should be) driven and focused at all times in supporting and meeting relevant platform and application SLAs.
- **Self-service and on-demand capacity:** This requires an online resource or customer portal that allows the customers to have complete visibility and awareness

of the virtual IaaS environment they currently utilize. It additionally allows customers to acquire, remove, manage, and report on resources, without the need to engage or speak with resources internally or with the provider.

- **High reliability and resilience:** To be effective, the requirement for automated distribution across the virtualized infrastructure is increasing and affording resilience, while enforcing and meeting SLA requirements.

IaaS Key Benefits

IaaS has a number of key benefits for organizations, which include but are not limited to the following:

- Usage metered and priced on the basis of units (or instances) consumed. This can also be billed back to specific departments or functions.
- The ability to scale up and down infrastructure services based on actual usage. This is particularly useful and beneficial when there are significant spikes and dips within the usage curve for infrastructure.
- Reduced cost of ownership. There is no need to buy assets for everyday use, no loss of asset value over time, and reduced costs of maintenance and support.
- Reduced energy and cooling costs along with "green IT" environment effect with optimum use of IT resources and systems.

Significant and notable providers in the IaaS space include Amazon, AT&T, Rackspace, Verizon/Terremark, and HP, among others.

PaaS

According to "The NIST Definition of Cloud Computing," in 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. The consumer does not manage or control the underlying cloud infrastructure, including network, servers, OSs, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment."[5]

PaaS and the cloud platform components have revolutionized the manner in which development and software has been delivered to customers and users over the past few years. The barrier for entry in terms of costs, resources, capabilities, and ease of use have dramatically reduced time to market—promoting and harvesting the innovative culture within many organizations.

PaaS Key Capabilities and Characteristics

Outside of the key benefits, PaaS should have the following key capabilities and characteristics:

- **Support multiple languages and frameworks:** PaaS should support multiple programming languages and frameworks, thus enabling the developers to code in whichever language they prefer or whatever the design requirements specify.

  In recent times, significant strides and efforts have been taken to ensure that open source stacks are both supported and utilized, thus reducing lock-in or issues with interoperability when changing CSPs.

- **Multiple hosting environments:** The ability to support a wide choice and variety of underlying hosting environments for the platform is key to meeting customer requirements and demands. Whether public cloud, private cloud, local hypervisor, or bare metal, supporting multiple hosting environments allows the application developer or administrator to migrate the application when and as required. This can also be used as a form of contingency and continuity and to ensure ongoing availability.

- **Flexibility:** Traditionally, platform providers provided features and requirements that they felt suited the client requirements, along with what suited their service offering and positioned them as the provider of choice, with limited options for the customers to move easily.

  This has changed drastically, with extensibility and flexibility now offered to meet the needs and requirements of developer audiences. This has been heavily influenced by open source, which allows relevant plug-ins to be quickly and efficiently introduced into the platform.

- **Allow choice and reduce lock-in:** Learning from previous horror stories and restrictions, proprietary meant red tape, barriers, and restrictions on what developers could do when it came to migration or adding features and components to the platform. Although the requirement to code to specific application programming interfaces (APIs) was made available by the provider, developers could run their apps in various environments based on commonality and standard API structures, ensuring a level of consistency and quality for customers and users.

- **Ability to auto-scale:** This enables the application to seamlessly scale up and down as required to accommodate the cyclical demands of users. The platform will allocate resources and assign these to the application, as required. This serves as a key driver for any seasonal organizations that experience spikes and drops in usage.

PaaS Key Benefits

PaaS has a number of key benefits for developers, which include but are not limited to these:

- OSs can be changed and upgraded frequently, including associated features and system services.
- Globally distributed development teams are able to work together on software development projects within the same environment.
- Services are available and can be obtained from diverse sources that cross national and international boundaries.
- Upfront and recurring or ongoing costs can be significantly reduced by utilizing a single vendor instead of maintaining multiple hardware facilities and environments.

Significant and notable providers in the PaaS space include Microsoft, OpenStack, and Google, among others.

SaaS

According to "The NIST Definition of Cloud Computing," in 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. The consumer does not manage or control the underlying cloud infrastructure including networks, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings."[6]

SaaS Delivery Models

Within SaaS, two delivery models are currently used:

- **Hosted application management (hosted AM):** The provider hosts commercially available software for customers and delivers it over the Web (Internet).
- **Software on demand:** The CSP gives customers network-based access to a single copy of an application created specifically for SaaS distribution (typically within the same network segment).

SaaS Benefits

Cloud computing provides significant and potentially limitless possibilities for organizations to run programs and applications that may previously have not been practical or feasible given the limitations of their own systems, infrastructure, or resources.

When utilizing and deploying the right middleware and associated components, the ability to run and execute programs with the flexibility, scalability, and on-demand self-service capabilities can present massive incentives and benefits for scalability, usability, reliability, productivity, and cost savings.

Clients can access their applications and data from anywhere at any time. They can access the cloud computing system using any computer linked to the Internet. Other capabilities and benefits related to the application include these:

- **Overall reduction of costs:** Cloud deployments reduce the need for advanced hardware to be deployed on the client side. Essentially, requirements to purchase high specification systems, redundancy, storage, and so on, to support applications are no longer necessary. From a customer perspective, a device to connect to the relevant application with the appropriate middleware is all that should be required.
- **Application and software licensing:** Customers no longer need to purchase licenses, support, and associated costs because licensing is leased and is relevant only when in use (covered by the provider). Additionally, purchasing of bulk licensing and the associated CapEx is removed and replaced by a pay-per-use licensing model.
- **Reduced support costs:** Customers save money on support issues because the relevant CSP handles them. Appropriately managed, owned, and operated streamlined hardware would, in theory, have fewer problems than a network of heterogeneous machines and OSs.

SaaS has a number of key benefits for organizations, which include but are not limited to these:

- Ease of use and limited administration.
- Automatic updates and patch management. The user is always running the latest version and most up-to-date deployment of the software release as well as any relevant security updates (no manual patching required).
- Standardization and compatibility. All users have the same version of the software release.
- Global accessibility.

Significant and notable providers in the SaaS space include Microsoft, Google, Salesforce.com, Oracle, and SAP, among others.

CLOUD DEPLOYMENT MODELS

Cloud deployment models fall into four main types of clouds: public, private, hybrid, and community.

Now that you are equipped with an understanding and appreciation of the cloud service types, you will learn how these services are merged into the relevant deployment models. The selection of a cloud deployment model will depend on any number of factors and may be heavily influenced by your organization's risk appetite, cost, compliance and regulatory requirements, and legal obligations, along with other internal business decisions and strategy.

The Public Cloud Model

According to NIST, "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."[7]

Public Cloud Benefits

The following are typical key drivers or benefits of a public cloud:

- Easy and inexpensive setup because the provider covers hardware, application, and bandwidth costs
- Streamlined and easy-to-provision resources
- Scalability to meet customer needs
- No wasted resources—pay as you consume

Given the increasing demands for public cloud services, many providers are now offering and remodeling their services as public cloud offerings. Significant and notable providers in the public cloud space include Amazon, Microsoft, Salesforce, and Google, among others.

The Private Cloud Model

According to NIST, "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."[8]

A private cloud is typically managed by the organization it serves; however, outsourcing the general management of this to trusted third parties may also be an option. A private cloud is typically available only to the entity or organization, its employees, contractors, and selected third parties.

The private cloud is also sometimes referred to as the *internal* or *organizational cloud*.

Private Cloud Benefits

Key drivers or benefits of a private cloud typically include these:

- Increased control over data, underlying systems, and applications
- Ownership and retention of governance controls
- Assurance over data location and removal of multiple jurisdiction legal and compliance requirements

Private clouds are typically more popular among large, complex organizations with legacy systems and heavily customized environments. Additionally, where significant technology investment has been made, it may be more financially viable to utilize and incorporate these investments within a private cloud environment than to discard or retire such devices.

The Hybrid Cloud Model

According to NIST, "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)."[9]

Hybrid cloud computing is gaining in popularity because it enables organizations to retain control of their IT environments, coupled with the convenience of allowing organizations to use public cloud service to fulfill non-mission-critical workloads and taking advantage of flexibility, scalability, and cost savings.

Hybrid Cloud Benefits

Key drivers or benefits of hybrid cloud deployments include these:

- Retain ownership and oversight of critical tasks and processes related to technology.
- Reuse previous investments in technology within the organization.
- Control the most critical business components and systems.
- Act as a cost-effective means of fulfilling noncritical business functions (utilizing public cloud components).

- Enhance cloud bursting and DR by hybrid cloud deployments. Cloud bursting allows for public cloud resources to be utilized when a private cloud workload has reached maximum capacity.

The Community Cloud Model

According to NIST, "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."[10]

Community clouds can be on-premises or offsite and should give the benefits of a public cloud deployment, while providing heightened levels of privacy, security, and regulatory compliance.

CLOUD CROSS-CUTTING ASPECTS

The deployment of cloud solutions, by its nature, is often deemed a technology decision; however, it's truly a business alignment decision. Although cloud computing no doubt enables technology to be delivered and utilized in a unique manner, potentially unleashing multiple benefits, the choice to deploy and consume cloud services should be a business decision, taken in line with the business or organization's overall strategy.

Why is it a business decision, you ask? Two distinct reasons:

- All technology decisions should be made with the overall business direction and strategy at the core.
- When it comes to funding and creating opportunities, these should be made at a business level.

A cloud transition's ability to directly support organizational business or mission goals and to express that message in a business manner is the difference between a successful project and a failed project in the eyes of the organization.

Architecture Overview

The architect is a planner, strategist, and consultant who sees the "big picture" of the organization. He understands current needs, thinks strategically, and plans long into the future. Perhaps the most important role of the architect today is to understand the business and how to design the systems that the business will require. This allows the architect to determine which system types, development, and configurations meet the identified business requirements while addressing any security concerns.

Enterprise security architecture provides the conceptual design of network security infrastructure and related security mechanisms, policies, and procedures. It links components of the security infrastructure as a cohesive unit with the goal of protecting corporate information. The Cloud Security Alliance (CSA) provides a general enterprise architecture (*Figure 1.4*). The CSA Enterprise Architecture is located at `https://cloudsecurityalliance.org/`.

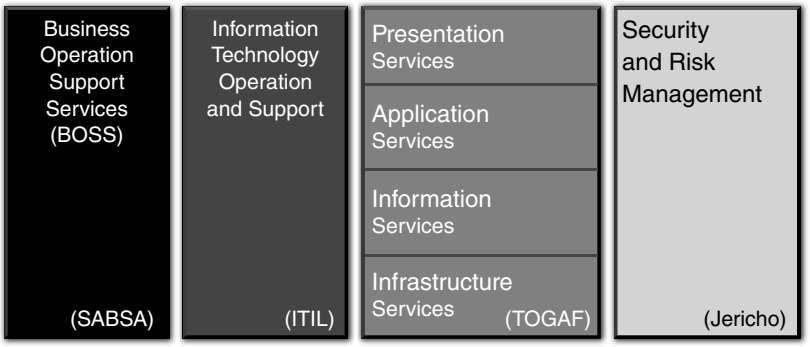

FIGURE 1.4 CSA Enterprise Architecture.

See the following sections for a starting point to reference the building blocks of the CSA Enterprise Architecture.

Sherwood Applied Business Security Architecture

Sherwood Applied Business Security Architecture (SABSA)[11] includes the following components, which can be used separately or together:

- Business Requirements Engineering Framework
- Risk and Opportunity Management Framework
- Policy Architecture Framework
- Security Services-Oriented Architecture Framework
- Governance Framework
- Security Domain Framework
- Through-Life Security Service Management and Performance Management Framework

Information Technology Infrastructure Library

Information Technology Infrastructure Library (ITIL)[12] is a group of documents that are used in implementing a framework for IT service management. ITIL forms a customizable framework that defines how service management is applied throughout an

organization. ITIL is organized into a series of five volumes: Service Strategy, Service Design, Service Transition, Service Operation, and Continual Service Improvement.

The Open Group Architecture Framework

The Open Group Architecture Framework (TOGAF)[13] is one of many frameworks available to the cloud security professional for developing an enterprise architecture. TOGAF provides a standardized approach that can be used to address business needs by providing a common lexicon for business communication. TOGAF is based on open methods and approaches to enterprise architecture, allowing the business to avoid a lock-in scenario from the use of proprietary approaches. TOGAF also provides for the ability to quantifiably measure return on investment (ROI) so that the business can use resources more efficiently.

Jericho/Open Group

The Jericho forum now is part of the Open Group Security Forum.[14] You can find the Jericho Forum Cloud Cube Model at `https://collaboration.opengroup.org/jericho/cloud_cube_model_v1.0.pdf`

Key Principles of an Enterprise Architecture

The following principles should be adhered to at all times:

- Define protections that enable trust in the cloud.
- Develop cross-platform capabilities and patterns for proprietary and open source providers.
- Facilitate trusted and efficient access, administration, and resiliency to the customer or consumer.
- Provide direction to secure information that is protected by regulations.
- Facilitate proper and efficient identification, authentication, authorization, administration, and auditability.
- Centralize security policy, maintenance operation, and oversight functions.
- Make access to information both secure and easy to obtain.
- Delegate or federate access control where appropriate.
- Ensure ease of adoption and consumption, supporting the design of security patterns.
- Make the architecture elastic, flexible, and resilient, supporting multitenant, multilandlord platforms.
- Ensure the architecture addresses and supports multiple levels of protection, including network, OS, and application security needs.

The NIST Cloud Technology Roadmap

The NIST Cloud Technology Roadmap helps CSPs develop industry-recommended, secure, and interoperable identity, access, and compliance management configurations and practices. It offers guidance and recommendations for enabling 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 CSPs are in terms of security capabilities and to plan a roadmap to meet the security needs of their business.[15]

There are a number of key components that the cloud security professional should comprehensively review and understand to determine which controls and techniques may be required to adequately address the requirements discussed in the following sections.

Interoperability

Interoperability defines how easy it is to move and reuse application components regardless of the provider, platform, OS, infrastructure, location, storage, format of data or APIs, how well applications work together, and how well new applications work with other solutions present in the business, organization, or provider's existing architecture.

Standards-based products, processes, and services are essential for entities to ensure the following:

- Investments do not become prematurely technologically obsolete.
- Organizations are able to easily change CSPs to flexibly and cost effectively support their mission.
- Organizations can economically acquire commercial and develop private clouds using standards-based products, processes, and services.

Interoperability mandates that those components should be replaceable by new or different components from different providers and continue to work, as should the exchange of data between systems.

Portability

Portability is a key aspect to consider when selecting CSPs because it can both help prevent vendor lock-in and deliver business benefits by allowing identical cloud deployments to occur in different CSP solutions, either for the purposes of DR or for the global deployment of a distributed single solution.

Availability

Systems and resource availability defines the success or failure of a cloud-based service. As a single point of failure (SPOF) for cloud-based services, where the service or cloud deployment loses availability, the customer is unable to access target assets or resources, resulting in downtime.

In many cases, CSPs are required to provide upward of 99.9 percent availability as per the SLA. Failure to do so can result in penalties, reimbursement of fees, loss of customers, loss of confidence, and ultimately brand and reputational damage.

Security

For many customers and potential cloud users, security remains the biggest concern, with security continuing to act as a barrier preventing them from engaging with cloud services.

As with any successful security program, the ability to measure, obtain assurance, and integrate contractual obligations to minimum levels of security are the keys to success.

Many CSPs now list their typical or minimum levels of security but will not list or publicly state specific security controls for fear of being targeted by attackers who would have the knowledge necessary to successfully compromise their networks.

Where such contracts and engagements require specific security controls and techniques to be applied, these are typically seen as extras. They incur additional costs and require that the relevant nondisclosure agreements (NDAs) be completed before engaging in active discussions.

In many cases, for smaller organizations, a move to cloud-based services significantly enhances their security controls, given that they may not have access to or possess the relevant security capabilities of a large-scale cloud computing provider.

The general rule of thumb for security controls and requirements in cloud-based environments is based on "if you want additional security, additional cost will be incurred." You can have almost whatever you want when it comes to cloud security—just as long as you can find the right provider and you are willing to pay for it.

Privacy

In the world of cloud computing, privacy presents a major challenge for both customers and providers alike. The reason for this is simple: no uniform or international privacy directives, laws, regulations, or controls exist, leading to a separate, disparate, and segmented mesh of laws and regulations being applicable depending on the geographic location where the information may reside (data at rest) or be transmitted (data in transit).

Although many of the leading providers of cloud services make provisions to ensure the location and legislative requirements (including contractual obligations) are met, this should never be taken as a given and should be specified within relevant SLAs and contracts. Given the true global nature and various international locations of cloud computing data centers, the potential for data to reside in two, three, or more locations around the world at any given time is a real possibility.

For many European entities and organizations, failure to ensure appropriate provisions and controls have been applied can violate EU data protection laws and obligations that can lead to various issues and implications.

Within Europe, privacy is seen as a human right and as such should be treated with the utmost respect. Not bypassing the various state laws across the United States and other geographic locations can make the job of the cloud architect extremely complex, requiring an intricate level of knowledge and controls to ensure that no such violations or breaches of privacy and data protection occur.

Resiliency

Cloud resiliency represents the ability of a cloud services data center and its associated components, including servers, storage, and so on, to continue operating in the event of a disruption, which may be equipment failure, power outage, or a natural disaster.

Given that most CSPs have a significantly higher number of devices and redundancy in place than a standard in-house IT team, resiliency should typically be far higher, with equipment and capabilities being ready to fail over, multiple layers of redundancy, and enhanced exercises to test such capabilities.

Performance

Cloud computing and high performance should go hand in hand at all times. Let's face it—if the performance is poor, you may not be a customer for very long. For optimum performance to be experienced through the use of cloud services, the provisioning, elasticity, and other associated components should always focus on performance.

The speed at which you can travel by boat depends on the engine and the boat design. The same applies for performance, which at all times should be focused on the network, the computer, the storage, and the data.

With these four elements influencing the design, integration, and development activities, performance should be boosted and enhanced throughout. It is always harder to refine and amend performance once design and development have been completed.

Governance

The term *governance* relating to processes and decisions looks to define actions, assign responsibilities, and verify performance. The same can be said and adopted for cloud

services and environments, where the goal is to secure applications and data when in transit and at rest. In many cases, cloud governance is an extension of the existing organizational or traditional business process governance, with a slightly altered risk and controls landscape.

Although governance is required from the commencement of a cloud strategy or cloud migration roadmap, it is seen as a recurring activity and should be performed on an ongoing basis.

A key benefit of many cloud-based services is the ability to access relevant reporting, metrics, and up-to-date statistics related to usage, actions, activities, downtime, outages, updates, and so on. This may enhance and streamline governance and oversight activities with the addition of scheduled and automated reporting.

Note that processes, procedures, and activities may require revision postmigration or movement to a cloud-based environment. Not all processes remain the same, with segregation of duties, reporting, and incident management forming a sample of processes that may require revision after the cloud migration.

SLAs

Think of a rulebook and legal contract all rolled into one document—that's what you have in terms of an SLA. In the SLA, the minimum levels of service, availability, security, controls, processes, communications, support, and many other crucial business elements are stated and agreed upon by both parties.

Many may argue that the SLAs are heavily weighted in favor of the CSP, but there are several key benefits when compared with traditional-based environments or in-house IT. These include downtime, upgrades, updates, patching, vulnerability testing, application coding, test and development, support, and release management. Many of these require the provider to take these areas and activities seriously; failing to do so affects their bottom line.

Note that not all SLAs cover the areas or focus points with which you may have issues or concerns. When this is not the case, every effort should be made to obtain clarity prior to engaging with the CSP services.

Auditability

Auditability allows for users and the organization to access, report, and obtain evidence of actions, controls, and processes that were performed or run by a specified user.

Similar to standard audit trails and systems logging, systems auditing and reporting are offered as standard by many of the leading CSPs.

From a customer perspective, increased confidence and the ability to have evidence to support audits, reviews, or assessments of object-level or systems-level access form key drivers.

From a stakeholder, management, and assessment perspective, auditability provides mechanisms to review, assess, and report user and systems activities. Auditability in non-cloud environments can focus on financial reporting, whereas cloud-based auditability focuses on actions and activities of users and systems.

Regulatory Compliance

Regulatory compliance is an organization's requirement to adhere to relevant laws, regulations, guidelines, and specifications relevant to its business, specifically dictated by the nature, operations, and functions it provides or utilizes to its customers. When the organization fails to meet or violates regulatory compliance regulations, punishment can include legal actions, fines, and, in limited cases, halting business operations or practices.

Key regulatory areas that are often included in cloud-based environments include but are not limited to the Payment Card Industry Data Security Standard (PCI DSS), the Health Insurance Portability and Accountability Act (HIPAA), the Federal Information Security Management Act (FISMA), and the Sarbanes-Oxley Act (SOX).

NETWORK SECURITY AND PERIMETER

Network security looks to cover all relevant security components of the underlying physical environment and the logical security controls that are inherent in the service or available to be consumed as a service (SaaS, PaaS, and IaaS). Two key elements need to be drawn out at this point:

- Physical environment security ensures that access to the cloud service is adequately distributed, monitored, and protected by underlying physical resources within which the service is built.
- Logical network security controls consist of link, protocol, and application layer services.

As a cloud customer and a CSP, both data and systems security are of utmost importance. The goal from both sides is to ensure the ongoing availability, integrity, and confidentiality of all systems and resources. Failure to do so has a negative impact from a customer, confidence, brand awareness, and overall security posture standpoint.

Taking into account that cloud computing requires a high volume of constant connections to and from the network devices, the always on and always available elements are necessary and essential.

In the cloud environments, the classic definition of a network perimeter takes on different meanings under different guises and deployment models.

- For many cloud networks, the perimeter is clearly the demarcation point.
- For other cloud networks, the perimeter transforms into a series of highly dynamic micro-borders around individual customer solutions or services (to the level of certain data sets and flows within a solution) within the same cloud, consisting of virtual network components.
- In other cloud networks, there is no clear perimeter at all. Although the network may be typically viewed as a perimeter and a number of devices within those perimeters communicating both internally and externally, this may be somewhat less clear and segregated in cloud computing networks.

Next, you will look at some of the add-on components that strengthen and enhance the overall security posture of cloud-based networks. You will see how to utilize them and learn why they play a fundamental function in technology deployments today.

CRYPTOGRAPHY

The need for the use of cryptography and encryption is universal for the provisioning and protection of confidentiality services in the enterprise. In support of that goal, the CCSP should ensure that he understands how to deploy and use cryptography services in a cloud environment. In addition, it's important to integrate strong key management services and a secure key management lifecycle into the cryptography solution.

Encryption

The need for confidentiality along with the requirement to apply additional security controls and mechanisms to protect information and communications is great. Whether it is encryption to a military standard or simply the use of self-signed certificates, everyone has different requirements and definitions of what a secure communications and cryptography-based infrastructure looks like. As with many areas of security, encryption can be subjective when you drill down into the algorithms, strengths, ciphers, implementation methods, and so on.

As a general rule of thumb, encryption mechanisms should be selected based on the information and data they protect, while taking into account requirements for access and general functions. The critical success factor for encryption is to enable secure and legitimate access to resources, while protecting and enforcing controls against unauthorized access.

The cloud architect and administrator should explore the appropriate encryption and access measures to ensure that proper separation of tenants' information and access is deployed within public cloud environments. Additionally, encryption and relevant controls need to be applied to private and hybrid cloud deployments to adequately and sufficiently protect communications between hosts and services across various network components and systems.

Data in Transit (Data in Motion)

Also described or termed *data in motion*, data in transit focuses on information or data while in transmission across systems and components typically across internal and external (untrusted) networks. Where information is crossing or traversing trusted and untrusted networks, the opportunity for interception, sniffing, or unauthorized access is heightened.

Data in transit can include the following scenarios:

- Data transiting from an end user endpoint (laptop, desktop, smart device, and so on) on the Internet to a web-facing service in the cloud
- Data moving between machines within the cloud (including between different cloud services), such as between a web virtual machine (VM) and a database
- Data traversing trusted and untrusted networks (cloud- and non-cloud-based environments)

Typically, the cloud architect is responsible for reviewing the way data in transit will be protected or secured at the design phase. Special consideration should be focused on how the cloud will integrate, communicate, and allow for interoperability across boundaries and hybrid technologies. Once implemented, the ongoing management and responsibility of data in transit resides in the correct application of security controls, including the relevant cryptography processes to handle key management.

Perhaps the best-known use of cryptography for the data in transit scenario is secure sockets layer (SSL) and transport layer security (TLS). TLS provides a transport layer–encrypted "tunnel" between email servers or message transfer agents (MTAs), whereas SSL certificates encrypt private communications over the Internet using private and public keys.

These cryptographic protocols have been in use for many years in the form of hypertext transfer protocol secure (HTTPS), typically to provide communication security over the Internet, but it has now become the standard and de facto encryption approach for browser-to-web host and host-to-host communications in both cloud and noncloud environments.

Recent increases show a number of cloud-based providers using multiple factors of encryption, coupled with the ability for users to encrypt their own data at rest within the cloud environment. The use of symmetric cryptography for key exchange followed by symmetric encryption for content confidentiality is also increasing.

This approach looks to bolster and enhance standard encryption levels and strengths of encryption. Additionally, IP security (IPSec), which has been used extensively, is a transit encryption protocol widely used and adopted for virtual private network (VPN) tunnels; it makes use of cryptography algorithms such as Triple DES (3DES) and Advanced Encryption Standard (AES).

Data at Rest

Data at rest focuses on information or data while stagnant or at rest (typically not in use) within systems, networks, or storage volumes. When data is at rest, appropriate and suitable security controls need to be applied to ensure the ongoing confidentiality and integrity of information.

Encryption of stored data, or data at rest, continues to gain traction for both cloud-based and non-cloud-based environments. The cloud architect is typically responsible for the design and assessment of encryption algorithms for use within cloud environments. Of key importance for both security and performance is the deployment and implementation of encryption on the target hosts and platforms.

The selection and testing of encryption form an essential component prior to ensuring performance impacts. In some cases, encryption can affect performance.

User interface (UI) response times and processor capabilities are up to a quarter or even half of the processor in an unencrypted environment. This varies depending on the type, strength, and algorithm. In high-performing environments with significant processor and utilization requirements, encryption of data at rest may not be included or utilized as standard.

Encryption of data at rest provides, assists, and assures organizations that opportunities for unauthorized access or viewing of data through information spills or residual data are further reduced.

Note that when information is encrypted on the CSP side and in the event of discrepancies or disputes with the providers, it may prove challenging to obtain or extract your data.

Key Management

In the old traditional banking environments, two people with keys were required to open the safe; this led to a reduced number of thefts, crimes, and bank robberies. Encryption, as with bank processes, should never be handled or addressed by a single person.

Encryption and segregation of duties should always go hand in hand. Key management should be separated from the provider hosting the data, and the data owners should be positioned to make decisions (these may be in line with organizational policies) but ultimately should be in a position to apply encryption, control, and manage key management processes, select the storage location for the encryption keys (on-premises in an isolated location is typically the best security option), and retain ownership and responsibilities for key management.

The Importance of Key Management

From a security perspective, you remove the dependency or assumption that the CSP is handling the encryption processes and controls correctly.

Also, you are not bound or restricted by shared keys or data spillage within the cloud environments because you have a unique and separate encryption mechanism to apply an additional level of security and confidentiality at a data and transport level.

Common Approaches to Key Management

For cloud computing key management services, the following two approaches are most commonly utilized:

- **Remote Key Management Service (KMS):** This is where the customer maintains the KMS on-premises. Ideally, the customer will own, operate, and maintain the KMS. This way the customer can control the information confidentiality, and the CSP can focus on the hosting, processing, and availability of services.

  Note that hybrid connectivity is required between the CSP and the cloud customer for the encryption and decryption to function.

- **Client-Side Key Management:** Similarly to the remote KMS approach, the client-side approach looks to put the customer or cloud user in complete control of the encryption and decryption keys.

  The main difference here is that most of the processing and control is done on the customer side. The CSP provides the KMS; however, the KMS resides on the customer's premises, where the customer generates, holds, and retains the keys. Note that this approach is typically utilized for SaaS environments and cloud deployments.

As with most areas of technology, access control is merging and aligning with other combined activities. Some of these are automated using single sign-on capabilities; others operate in a standalone, segregated fashion.

The combination of access control and effective management of those technologies, processes, and controls has given rise to identity and access management (IAM). In a nutshell, IAM includes people, processes, and systems that manage access to enterprise resources. This is achieved by ensuring that the identity of an entity is verified (who are they, can they prove who they are) and then granting the correct level of access based on the assets, services, and protected resources being accessed.

IAM typically looks to utilize a minimum of two—preferably three or more—factors of authentication. Within cloud environments, services should include strong authentication mechanisms for validating users' identities and credentials. In line with best practice, one-time passwords should be utilized as a risk reduction and mitigation technique.

The key phases that form the basis and foundation for IAM in the enterprise include the following:

- Provisioning and deprovisioning
- Centralized directory services
- Privileged user management
- Authentication and access management

Each is discussed in the following sections.

Provisioning and Deprovisioning

Provisioning and deprovisioning are critical aspects of access management. Think of setting up and removing users. In the same way as you would set up an account for a user entering your organization requiring access to resources, provisioning is the process of creating accounts to allow users to access appropriate systems and resources within the cloud environment.

The ultimate goal of user provisioning is to standardize, streamline, and create an efficient account creation process, while creating a consistent, measurable, traceable, and auditable framework for providing access to end users.

Deprovisioning is the process whereby a user account is disabled when the user no longer requires access to the cloud-based services and resources. This is not just limited to a user leaving the organization but may also be due to a user changing a role, function, or department.

Deprovisioning is a risk-mitigation technique to ensure that authorization creep or additional and historical privileges are not retained, thus granting access to data, assets, and resources that are not necessary to fulfill the job role.

Centralized Directory Services

As when building a house or large structure, the foundation is key. In the world of IAM, the directory service forms the foundation for IAM and security both in an enterprise environment and within a cloud deployment. A directory service stores, processes, and facilitates a structured repository of information stored, coupled with unique identifiers and locations.

The primary protocol in relation to centralized directory services is Lightweight Directory Access Protocol (LDAP), built and focused on the X.500 standard.[16] LDAP works as an application protocol for querying and modifying items in directory service providers like Active Directory. Active Directory is a database-based system that offers authentication, directory, policy, and other services to a network.

Essentially, LDAP acts as a communication protocol to interact with Active Directory. LDAP directory servers store their data hierarchically (similar to domain name system [DNS] trees and UNIX file structures) with a directory record's distinguished name (DN) read from the individual entries back through the tree, up to the top level.

Each entry in an LDAP directory server is identified through a DN access to directory services, should be part of the IAM solution, and should be as robust as the core authentication modes used.

The use of privileged identity management (PIM) features is strongly encouraged for managing access of the administrators of the directory. If these are hosted locally rather than in the cloud, the IAM service requires connectivity to the local LDAP servers, in addition to any applications and services for which it is managing access.

Within cloud environments, directory services are heavily utilized and depended upon as the go-to trusted source by the IAM framework as a security repository of identity and access information. The same can be said for federated environments. Again, trust and confidence in the accuracy and integrity of the directory services are must-haves.

Privileged User Management

As the names implies, privileged user management focuses on the process and ongoing requirements to manage the lifecycle of user accounts with highest privileges in a system. Privileged accounts typically carry the highest risk and impact because compromised privileged user accounts can lead to significant permissions and access rights being obtained, thus allowing the user or attacker to access resources and assets that may negatively affect the organization.

The key components from a security perspective relating to privileged user management should, at a minimum, include the ability to track usage, authentication successes and failures, and authorization times and dates; log successful and failed events; enforce password management; and contain sufficient levels of auditing and reporting related to privileged user accounts.

Many organizations monitor this level of information for standard or general users, which would be beneficial and useful in the event of an investigation; however, the privileged accounts should capture this level of detail by default because attackers often target and compromise a general or standard user, with the view to escalating privileges to a more privileged or admin account. Not forgetting that a number of these components are technical by nature, the overall requirements that are used to manage these should be driven by organizational policies and procedures.

Note that segregation of duties can form an extremely effective mitigation and risk-reduction technique around privileged users and their ability to effect major changes.

Authorization and Access Management

Access to devices, systems, and resources forms a key driver for use of cloud services (broad network access); without it, the overall benefits that the service may provide are reduced to the enterprise, and legitimate business or organizational users are isolated from their resources and assets.

In the same way that users require authorization and access management to be operating and functioning to access the required resources, security requires these service components to be functional, operational, and trusted to enforce security within cloud environments.

In its simplest form, authorization determines the user's right to access a certain resource. (Think of entry onto a plane with your reserved seat or when you may be visiting an official residence or government agency to visit a specified person.)

Access management is focused on the manner and way in which users can access relevant resources, based on their credentials and characteristics of their identity. (Think of a bank or highly secure venue—only certain employees or personnel can access the main safe or highly sensitive areas.)

Note that both authorization and access management are point-in-time activities that rely on the accuracy and ongoing availability of resources and functioning processes, segregation of duties, privileged user management, password management, and so on, to operate and provide the desired levels of security. If one of the mentioned activities is not carried out regularly as part of an ongoing managed process, it can weaken the overall security posture.

DATA AND MEDIA SANITIZATION

By its nature, cloud-based environments are typically hosting multiple types, structures, and components of data among various resources, components, and services for users to access. If you want to leave or migrate from one CSP to another, this may be possible with little hassle, although other entities have experienced significant challenges in removing and exporting their large amounts of structured data from one provider to another. This is where vendor lock-in and interoperability elements come to the fore. It's also necessary to consider data and media sanitization. The ability to safely remove all data from a system or media, rendering it inaccessible, is critical to ensuring confidentiality and to managing a secure lifecycle for data in the cloud.

Vendor Lock-In

Vendor lock-in highlights where a customer may be unable to leave, migrate, or transfer to an alternate provider because of technical or nontechnical constraints. Typically, this could be based on the technology, platforms, or system design that may be proprietary or because of a dispute between the provider and the customer. Vendor lock-in poses a real risk for an organization that may not be in a position to leave the current provider or indeed continue with business operations and services.

Additionally, where a specific proprietary service or structure has been used to store your vast amounts of information, this may not support the intelligent export into a structured format. For example, how many organizations would be pleased with 100,000 records being exported into a flat-based text file? Open APIs are being strongly championed as a mechanism to reduce this challenge.

Aside from the hassle and general issues associated with reconstructing and formatting large data sets into a format that could be imported and integrated into a new cloud service or CSP, the challenge related to secure deletion or the sanitization of digital media remains a largely unsolved issue among CSPs and cloud customers alike.

Most organizations have failed to assess or factor in this challenge in the absence of a cloud computing strategy, and ultimately many have not put highly sensitive or regulated data in cloud-based environments as yet. This is likely to change with the shift toward compliant clouds and cloud-based environments aligned with certification standards such as ISO 27001/2, SOC 2, and PCI DSS among other international frameworks.

In the absence of degaussing, which is not a practical or realistic option for cloud environments, the approach for rendering data unreadable should be the first option taken (assuming the physical destruction of storage areas is not feasible). Adopting a security mind-set, if you can restrict the availability, integrity, and confidentiality of the data, you can then make the information unreadable, which will act as the next best method to secure deletion. How might you achieve this in cloud-based environments?

Cryptographic Erasure

A fairly reliable way to sanitize a device is to erase or overwrite the data it contains. With the recent developments in storage devices, most now contain built-in sanitize commands that enable users and custodians to sanitize media in a simple and convenient format. Although these commands are mostly effective when implemented and initiated correctly, like all technological commands, it is essential to verify their effectiveness and accuracy.

Where possible (this may not apply to all cloud-based environments), erase each block, overwrite all with a known pattern, and erase them again.

When done correctly, a complete erasure of the storage media eliminates risks related to key recovery (where stored locally—yes, this is a common mistake), side-channel attacks on controller to recover information about the destroyed key, and future attacks on the cryptosystem.

Note that key destruction on its own is not a comprehensive approach because the key may be recovered using forensic techniques.

Data Overwriting

Although it is not inherently secure and does not make the data irretrievable, overwriting data multiple times can make the task of retrieval far more complex, challenging, and time consuming. This technique may not be sufficient if you are hosting highly sensitive, confidential, or regulated information within cloud deployments.

When you delete files and data, they become invisible to the user; however, the space they inhabit in the storage media is made available for other information and data to be written to by the system and storage components as part of normal usage of the storage media. The challenge and risk with this is that forensic investigators and relevant toolsets can retrieve this information in a matter of minutes, hours, or days.

Where possible, overwriting data multiple times helps extend the time and efforts required to retrieve the relevant information and may make the storage components or partitions unattractive to potential attackers or those focused on retrieving the information.

WARNING Given enough time, effort, and resources in the absence of degaussing media, these approaches may not be sufficient to evade a determined attacker or reviewer from retrieving relevant information. What it may do is dissuade or make the task too challenging for a novice, intermediate, or opportunist attacker, who could decide to target easier locations or storage mediums.

VIRTUALIZATION SECURITY

Virtualization technologies enable cloud computing to become a real and scalable service offering due to the savings, sharing, and allocations of resources across multiple tenants and environments. As with all enabling technologies, the specified deployment and manner in which the solution is deployed may allow attackers to target relevant components and functions with the view to obtain unauthorized access to data, systems, and resources.

In the world of cloud computing, virtualization represents one of the key targets for the attackers. Specifically, although virtualization may introduce technical vulnerabilities based on the solution, the single most critical component to enable the technology to function in the manner for which it was developed, along with enforcing the relevant technical and nontechnical security controls, is the hypervisor.

The Hypervisor

The role of the hypervisor is a simple one: to allow multiple OSs to share a single hardware host (with each OS appearing to have the host's processor, memory, and resources to itself).

Think of a management console. Effectively, this is what the hypervisor does—intelligently controlling the host processor and resources, prioritizing and allocating what is needed to each OS, while ensuring there are no crashes and the neighbors do not upset each other.

Now you will dig a little deeper, with the goal of learning the security elements associated with VMs.

- **Type 1 hypervisor:** There are many accounts, definitions, and versions of what the distinctions between Type 1 and Type 2 hypervisors are (and are not), but with the view to keeping it simple, this book refers to Type 1 hypervisors as those running directly on the hardware with VM (guest operating system) resources provided by the hypervisor.

  These are also referred to as bare metal hypervisors. Examples of these include VMware ESXi and Citrix XenServer.

- **Type 2 hypervisor:** Type 2 hypervisors run on a host OS to provide virtualization services. Examples of Type 2 are VMware Workstation and Virtual Box.

In summary, Type 1 relates to hardware, and Type 2 relates to an OS.

Security Types

From a security perspective, you'll now explore which of the hypervisors provides a more robust security posture and which is more targeted by attackers.

- **Type 1 security:** Type 1 hypervisors significantly reduce the attack surface over Type 2 hypervisors. Type 1 hypervisor vendors also control relevant software that comprise and form the hypervisor package, including the virtualization functions and OS functions, such as devices drivers and input/output (I/O) stacks.

  Because the vendors have control over the relevant packages, they can reduce the likelihood of malicious software being introduced into the hypervisor foundation and introducing or exposing the hypervisor layer.

  The limited access and strong control over the embedded OS greatly increase the reliability and robustness of Type 1 hypervisors.
- **Type 2 security:** Because Type 2 hypervisors are OS based, they are more attractive to attackers, given that there are far more vulnerabilities associated with the OS as well as other applications that reside within the OS layer.

  A lack of standardization on the OS and other layers can open up additional opportunities and exposures that might make the hypervisor susceptible to attack and compromise.

Where technology, hardware, and software standardization can be used effectively, this can significantly reduce the risk landscape and increase the security posture.

COMMON THREATS

Threats form a real and ever-evolving challenge for organizations to counteract and defend against. Whether they are cloud specific or general disruptions to business and technology, threats can cause significant issues, outages, poor performance, and catastrophic impacts should they materialize.

Many of the top risks identified in the research paper *The Notorious Nine: Cloud Computing Top Threats in 2013*, published by the Cloud Security Alliance's Top Threats Working Group, remain a challenge for non-cloud-based environments and organizations alike. What this illustrates is the consistent challenges faced by entities today, altered and amplified by different technology deployments, such as cloud computing.[17]

Data Breaches

Not new to the security practitioner and company leaders, this age-old challenge continues to dominate headlines and new stories around the world. Whether it is a lost laptop

that is unencrypted or side channel timing attacks on VMs, what cloud computing has done is widen the scope and coverage for data breaches.

Given the nature of cloud deployments and multitenancy, VMs, shared databases, application design, integration, APIs, cryptography deployments, key management, and multiple locations of data all combine to provide a highly amplified and dispersed attack surface, leading to greater opportunity for data breaches.

Cloud security professionals can expect to be facing far more data breaches and loss of organizational and personal information as the adoption of the cloud and further use of mobile devices continue to increase. This is in large measure due to the rise of smart devices, tablets, increased workforce mobility, bring your own device (BYOD), and other factors, such as the historical challenge of lost devices, compromised systems, and traditional forms of attacks, coupled with the previously listed factors related to the cloud.

Depending on the data and information classification types, any data breaches or suspected breaches of systems security controls may require mandatory breach reporting to relevant agencies, entities, or bodies. This can include healthcare information (HIPAA), personally identifiable information (Directive 95/46/EC of the European Parliament and of the Council on the protection of individuals with regard to the processing of personal data and on the free movement of such data), and credit card information (PCI DSS). Significant fines may be imposed on organizations that cannot illustrate sufficient duty of care or security controls being implemented to prevent such data breaches. These vary greatly depending on the industry, sector, geographic location, and nature of the information.

Data Loss

Not to be confused with a data breach, data loss refers to the loss of information, deletion, overwriting, corruption, or integrity related to the information stored, processed, or transmitted within cloud environments.

Data loss within cloud environments can present a significant threat and challenge to organizations. The reasons for this can be illustrated by the following questions:

- Does the provider or customer have responsibility for data backup?
- If backup media containing the data is obtained, does this include all data or only a portion of the information?
- Where data has become corrupt or overwritten, can an import or restore be performed?
- Where accidental data deletion has occurred from the customer side, will the provider facilitate the restoration of systems and information in multitenancy environments or on shared platforms?

Note that when the customer uploads encrypted information to the cloud environment, the encryption keys become a critical component to ensure data is not lost and remains available. The loss of the relevant encryption keys constitutes data loss because the information will no longer be available for use in the absence of the keys.

Security can from time to time come back to haunt you if it is not owned, operated, and maintained effectively and efficiently.

Account or Service Traffic Hijacking

This is not a cloud-specific threat but one that has been a constant thorn and challenge for relevant security professionals to combat through the years. Account and service traffic hijacking has long been targeted by attackers, using methods such as phishing, more recently smishing (SMS phishing), spear phishing (targeted phishing attacks), and exploitation of software and other application-related vulnerabilities.

The key component of these attack methods, when successful, allows for the attackers to monitor and eavesdrop on communications, sniff and track traffic, capture relevant credentials, and access and alter account and user profile characteristics (changing passwords and more).

Of late, attackers are utilizing compromised systems, accounts, and domains as a smokescreen to launch attacks against other organizations and entities, making the source of the attack appear to be from suppliers, third parties, competitors, or other legitimate organizations that have no knowledge or awareness of having been compromised.

Insecure Interfaces and APIs

For users to access cloud computing assets and resources, they utilize the APIs made available by the CSP. Key functions of the APIs, including the provisioning, management, and monitoring, are performed utilizing the provider interfaces. For the security controls and availability of resources to function in the way that they were designed, use of the provider APIs is required to prevent against deliberate and accidental attempts to circumvent policies and controls.

Sounds simple enough, right? In an ideal world, that may be true, but for the modern and evolving cloud landscape, that challenge is amplified with relevant third parties, organizations, and customers (depending on deployment) building additional interfaces and "bolt on" components to the API, which significantly increase the complexity, resulting in a multilayered API. This can result in credentials being passed to third parties or consumed insecurely across the API and relevant stack components.

Note that most providers make concerted efforts to ensure the security of their interfaces and APIs; however, any variations or additional components added on from the consumer or other providers can reduce the overall security posture and stance.

Denial of Service

By their nature, denial-of-service (DoS) attacks prevent users from accessing services and resources from a specified system or location. This can be done using any number of attack vectors available but typically look to target buffers, memory, network bandwidth, or processor power.

With cloud services relying ultimately on availability to service and enable connectivity to resources from customers, when DoS attacks are targeted at cloud environments, they can create significant challenges for the provider and customer alike.

Distributed denial-of-service (DDoS) attacks are launched from multiple locations against a single target. Work with the cloud security architect to ensure that system design and implementation do not create an SPOF that can expose an entire system to failure if a DoS or DDoS attack is successfully launched against a system.

Note that although it's widely touted by the media and feared by organizations worldwide, many believe that DoS attacks require large volumes of traffic to be successful. This is not always the case; asymmetric application-level payload attacks have measured success with as little at 100–150Kbps packets.

Malicious Insiders

When looking to secure the key assets of any organization, three primary components are essential—people, processes, and technology. People tend to present the single largest challenge to security due to the possibility of a disgruntled, rogue, or simply careless employee or contractor exposing sensitive data either by accident or on purpose.

According to CERT, malicious insider threats to an organization can come from "a current or former employee, contractor, or other business partner who has or had authorized access to an organization's network, system, or data and intentionally exceeded or misused that access in a manner that negatively affected the confidentiality, integrity, or availability of the organization's information or information systems."[18]

Abuse of Cloud Services

Think of the ability to have previously unobtainable and unaffordable computing resources available for a couple of dollars an hour. Well, that is exactly what cloud computing provides—an opportunity for businesses to have almost unlimited scalability and flexibility. The challenge for many organizations is that this scalability and flexibility are provided across the same platforms or resources that attackers can access and use to execute dictionary attacks, execute DoS attacks, crack encryption passwords, or host illegal software and materials for widespread distribution. Note that the power of the cloud is not always used in the manner for which it is offered to users.

Insufficient Due Diligence

Cloud computing has created a revolution among many users and companies with regard to how they utilize technology-based solutions and architectures. As with many such technology changes and revolutions, some have acted before giving the appropriate thought and due care to what a secure architecture would look like and what would be required to implement one.

Cloud computing has, for many organizations, become that rash decision—intentionally or unintentionally. The change in roles, focus, governance, auditing, reporting, strategy, and other operational elements requires a considerable investment on the part of the business in a thorough risk-review process, as well as amendments to business processes.

Given the immaturity of the cloud computing market, many entities and providers are still altering and refining the way they operate. There will be acquisitions, changes, amendments, and revisions in the way in which entities offer services, which can influence both customers and partners.

Finally, when the dust settles in the race for cloud space, pricing may vary significantly, rates and offerings may be reduced or inflated, and cyber attacks could force customers to review and revise their selection of a CSP. Should your provider go bankrupt, are you in a position to change CSPs in a timely and seamless manner?

It is incumbent upon the cloud security professional to ensure that both due care and due diligence are being exercised in the drive to the cloud.

- **Due diligence** is the act of investigating and understanding the risks a company faces.
- **Due care** is the development and implementation of policies and procedures to aid in protecting the company, its assets, and its people from threats.

Note that cloud companies may merge, be acquired, go bust, change services, and ultimately change their pricing model. Those that fail to carry out the appropriate due diligence activities may in fact be left with nowhere to go or turn to unless they introduce compensating controls to offset such risks (potentially resulting in less financial benefit).

Shared Technology Vulnerabilities

For CSPs to effectively and efficiently deliver their services in a scalable way, they share infrastructure, platforms, and applications among tenants and potentially with other providers. This can include the underlying components of the infrastructure, resulting in shared threats and vulnerabilities.

Where possible, providers should implement a layered approach to securing the various components. A defense-in-depth strategy should include compute, storage, network, application, and user security enforcement and monitoring. This should be universal, regardless of whether the service model is IaaS, PaaS, or SaaS.

SECURITY CONSIDERATIONS FOR DIFFERENT CLOUD CATEGORIES

Security can be a subjective issue, viewed differently across different industries, companies, and users, based on their needs, desires, and requirements. Many of these actions and security appetites are strongly influenced by compliance and other regulatory requirements.

IaaS Security

Within IaaS, a key emphasis and focus must be placed on the various layers and components stemming from the architecture to the virtual components. Given the reliance and focus placed on the widespread use of virtualization and the associated hypervisor components, this must be a key focus as an attack vector to gain access to or disrupt a cloud service.

The hypervisor acts as the abstraction layer that provides the management functions for required hardware resources among VMs.

- **VM attacks:** Cloud servers contain tens of VMs. These VMs may be active or offline and, regardless of state, are susceptible to attacks. Active VMs are vulnerable to all traditional attacks that can affect physical servers.

  Once a VM is compromised, VMs on the same physical server can attack each other because they share the same hardware and software resources, including memory, device drivers, storage, and hypervisor software.

- **Virtual network:** The virtual network contains the virtual switch software that controls the movement of traffic between the virtual network interface cards (NICs) of the installed VMs and the physical NICs of the host.

- **Hypervisor attacks:** Malicious hackers consider the hypervisor a potential target because of the greater control afforded by lower layers in the system. Compromising the hypervisor enables control over the installed VMs, the physical system, and the hosted applications.

  Common attacks include hyperjacking (installing a rogue hypervisor that can take complete control of a server), such as SubVir, Blue Pill (hypervisor rootkit using AMD secure virtual machine [SVM]), Vitriol (hypervisor rootkit using Intel VT-x), and direct kernel structure manipulation (DKSM).

  Another common attack is the VM escape, which is done by crashing the guest OS to get out of it and running an arbitrary code on the host OS. This allows malicious VMs to take complete control of the host OS.

- **VM-based rootkits (VMBRs):** These rootkits act by inserting a malicious hypervisor on the fly or modifying the installed hypervisor to gain control over the host workload.
- **Virtual switch attacks:** The virtual switch is vulnerable to a wide range of layer II attacks such as manipulation or modification of the virtual switch's configuration, VLANs and trust zones, and ARP tables.
- **DoS attacks:** DoS attacks in a virtual environment form a critical threat to VMs, along with all other dependent and associated services.

  Note that not all DoS attacks are from external attackers.

  These attacks can be the direct result of misconfigurations at the hypervisor, which allows a single VM instance to consume and utilize all available resources. In the same manner as a DoS attack renders resources unavailable to users attempting to access them, misconfigurations at the hypervisor restrict any other VM running on the same physical machine. This prevents network hosts from functioning appropriately because of the resources being consumed and utilized by a single device.

  Hypervisors prevent any VM from gaining 100 percent usage of any shared hardware resources, including CPU, RAM, network bandwidth, and other memory. Appropriately configured hypervisors detect instances of resource hogging and take appropriate actions, such as restarting the VM in an effort to stabilize or halt any processes that may be causing the abuse.
- **Colocation:** Multiple VMs residing on a single server and sharing the same resources increase the attack surface and the risk of VM-to-VM or VM-to-hypervisor compromise. On the other hand, when a physical server is off, it is safe from attacks. However, when a VM comes offline, it is still available as VM image files that are susceptible to malware infections and patching.

Provisioning tools and VM templates are exposed to different attacks that attempt to create new unauthorized VMs or patch the VM templates. This infects the other VMs that will be cloned from this template.

These new categories of security threats are a result of the new, complex, and dynamic nature of the cloud virtual infrastructure, as follows:

- **Multitenancy:** By design, different users within a cloud share the same applications and the physical hardware to run their VMs. As a result, information leakage as well as an increase in the attack surface and the risk of VM-to-VM or VM-to-hypervisor compromise can occur.

- **Loss of control:** Users are typically not aware of the location of their data and services, whereas the CSPs host and run VMs without being aware of their contents.
- **Network topology:** Cloud architecture is dynamic due to the fact that existing workloads change over time because of the creation and removal of VMs. In addition, the abilities of VMs to migrate from one host to another leads to the rise of nonpredefined network topologies.
- **Logical network segmentation:** Within IaaS, the requirement for isolation alongside the hypervisor remains a key and fundamental activity to reduce external sniffing, monitoring, and interception of communications and others within the relevant segments.

  When assessing relevant security configurations and connectivity models, VLANs, NATs, bridging, and segregation provide viable options to ensure the overall security posture remains strong, flexible, and constant, as opposed to other mitigation controls that may affect the overall performance.
- **No physical endpoints:** Due to server and network virtualization, the number of physical endpoints (such as switches, servers, and NICs) is reduced. These physical endpoints are traditionally used in defining, managing, and protecting IT assets.
- **Single point of access (SPOA) or SPOF:** Hosts have a limited number of access points (NICs) available to all VMs.

  This represents a critical security vulnerability: compromising these access points opens the door to compromise the VMs, the hypervisor, or the virtual switch.

The Cloud Security Alliance Common Controls Matrix (CCM) provides a good go-to guide for specific risks for SaaS, PaaS, and IaaS.[19]

PaaS Security

PaaS security involves four main areas, each of which is discussed in the following sections.

System and Resource Isolation

PaaS tenants should not have shell access to the servers running their instances (even when virtualized). The rationale behind this is to limit the chance and likelihood of configuration or system changes affecting multiple tenants. Where possible, administration facilities should be restricted to siloed containers to reduce this risk.

Careful consideration should be given before access is provided to the underlying infrastructure hosting a PaaS instance. In enterprises, this may have less to do with

malicious behavior and more to do with efficient cost control; it takes time and effort to undo tenant-related fixes to their environments.

User-Level Permissions

Each instance of a service should have its own notion of user-level entitlements (permissions). If the instances share common policies, appropriate countermeasures and controls should be enabled by the cloud security professional to reduce authorization creep or the inheritance of permissions over time.

However, it is not all a challenge; the effective implementation of distinct and common permissions can yield significant benefits when implemented across multiple applications within the cloud environment.

User Access Management

User access management enables users to access IT services, resources, data, and other assets. Access management helps to protect the availability, integrity, and confidentiality (AIC) of these assets and resources, ensuring that only those authorized to use or access these are permitted access.

In recent years, traditional standalone access control methods have become less utilized, with more holistic approaches to unify the authentication of users becoming favored. (This includes single sign-on.) For user access management processes and controls to function effectively, a key emphasis is placed on the agreement, implementation of the rules, and organizational policies for access to data and assets.

The key components of user access management include but are not limited to the following:

- **Intelligence:** Requires collection, analysis, auditing, and reporting against rule-based criteria, typically based on organizational policies.
- **Administration:** The ability to perform onboarding or changing account access on systems and applications.

  These solutions or toolsets should enable automation of tasks that were typically or historically performed by personnel within the operations or security function.
- **Authentication:** Provides assurance and verification in real time as to the user being who she claims to be, accompanied by relevant credentials (such as passwords).
- **Authorization:** Determines the level of access to grant each user based on policies, roles, rules, and attributes. The principle of least privilege should always be applied (that is, only what is specifically required to fulfill the job functions).

Note that User Access Management enables organizations to avail benefits across the areas of security, operational efficiencies, user administration, auditing, and reporting along with other onboarding components; however, it can be difficult to implement for historical components or environments.

Protection Against Malware, Backdoors, and Trojans

Traditionally, development and other teams create backdoors to enable administrative tasks to be performed.

The challenge with these is that once backdoors are created, they provide a constant vector for attackers to target and potentially gain access to the relevant PaaS resources. You have heard of the story in which attackers gained access through a backdoor, only to create additional backdoors while removing the legitimate backdoors, essentially holding the systems, resources, and associated services hostage.

More recently, attackers have utilized embedded and hardcoded malware as a method of obtaining unauthorized access and retaining this access for a prolonged and extended period. Most notably, malware has been placed in point-of-sale (PoS) devices, handheld card-processing devices, and other platforms, thereby divulging large amounts of sensitive data (including credit card numbers, customer details, and so on).

As with SaaS, web application and development reviews should go hand in hand. Code reviews and other software development lifecycle checks are essential to ensure that the likelihood of malware, backdoors, Trojans, and other potentially harmful vectors is reduced significantly.

SaaS Security

SaaS security involves three main areas, each of which is discussed in the following sections.

Data Segregation

Multitenancy is one of the major characteristics of cloud computing. As a result of multitenancy, multiple users can store their data using the applications that SaaS provides. Within these architectures, the data of various users will reside at the same location or across multiple locations and sites. With the appropriate permissions or using attack methods, the data of customers may become visible or possible to access.

Typically, in SaaS environments, this can be achieved by exploiting code vulnerabilities or injecting code within the SaaS application. If the application executes this code without verification, there is a high potential of success for the attacker to access or view other customers' or tenants' data.

A SaaS model should therefore ensure a clear segregation for each user's data. The segregation must be ensured not only at the physical level but also at the application level. The service should be intelligent enough to segregate the data from different users. A malicious user can use application vulnerabilities to hand-craft parameters that bypass security checks and access sensitive data of other tenants.

Data Access and Policies

When allowing and reviewing access to customer data, the key aspect to structuring a measurable and scalable approach begins with the correct identification, customization, implementation, and repeated assessments of the security policies for accessing data.

The challenge associated with this is to map existing security policies, processes, and standards to meet and match the policies that the CSP enforces. This may mean revising existing internal policies or adopting new practices whereby users can only access data and resources relevant to their job function and role.

The cloud must adhere to these security policies to avoid intrusion or unauthorized users viewing or accessing data.

The challenge from a CSP perspective is to offer a solution and service that is flexible enough to incorporate the specific organizational policies put forward by the organization, while also being positioned to provide a boundary and segregation among the multiple organizations and customers within a single cloud environment.

Web Application Security

Because SaaS resources are required to be always on and availability disruptions kept to a minimum, security vulnerabilities within the web application(s) carry significant risk and potential impact for the enterprise. Vulnerabilities, no matter what risk categorization, present challenges for CSPs and customers alike. Given the large volume of shared and colocated tenants within SaaS environments, if a vulnerability is exploited, both the cloud customer and the service provider may experience catastrophic consequences.

As with traditional web application technologies, cloud services rely on a robust, hardened, and regularly assessed web application to deliver services to its users. The fundamental difference with cloud-based services versus traditional web applications is their footprint and the attack surface they will present.

In the same way that web application security assessments and code reviews are performed on applications prior to release, this becomes even more crucial when dealing with cloud services. The failure to carry out web application security assessments and code reviews may result in unauthorized access, corruption, or other integrity issues affecting the data, along with a loss of availability.

Finally, web applications introduce new and specific security risks that may not be counteracted or defended against by traditional network security solutions (firewalls, intrusion detection systems [IDSs], intrusion prevention systems [IPSs], and so on). The nature and manner in which web application vulnerabilities and exploits operate may not be identified or may appear legitimate to the network security devices designed for non-cloud architectures.

OPEN WEB APPLICATION SECURITY PROJECT TOP TEN SECURITY THREATS

The Open Web Application Security Project (OWASP) has provided the 10 most critical web application security threats that should serve as a minimum level for application security assessments and testing.

The OWASP top 10 covers the following categories:

- "**A1—Injection:** Injection flaws, such as SQL, OS, and LDAP injection occur when untrusted data is sent to an interpreter as part of a command or query. The attacker's hostile data can trick the interpreter into executing unintended commands or accessing data without proper authorization.
- "**A2—Broken Authentication and Session Management:** Application functions related to authentication and session management are often not implemented correctly, allowing attackers to compromise passwords, keys, or session tokens, or to exploit other implementation flaws to assume other users' identities.
- "**A3—Cross-Site Scripting (XSS):** XSS flaws occur whenever an application takes untrusted data and sends it to a web browser without proper validation or escaping. XSS allows attackers to execute scripts in the victim's browser, which can hijack user sessions, deface websites, or redirect the user to malicious sites.
- "**A4—Insecure Direct Object References:** A direct object reference occurs when a developer exposes a reference to an internal implementation object, such as a file, directory, or database key. Without an access control check or other protection, attackers can manipulate these references to access unauthorized data.
- "**A5—Security Misconfiguration:** Good security requires having a secure configuration defined and deployed for the application, frameworks, application server, web server, database server, and platform. Secure settings should be defined, implemented, and maintained, as defaults are often insecure. Additionally, software should be kept up to date.

- "**A6—Sensitive Data Exposure:** Many web applications do not properly protect sensitive data, such as credit cards, tax IDs, and authentication credentials. Attackers may steal or modify such weakly protected data to conduct credit card fraud, identity theft, or other crimes. Sensitive data deserves extra protection such as encryption at rest or in transit, as well as special precautions when exchanged with the browser.
- "**A7—Missing Function Level Access Control:** Most web applications verify function-level access rights before making that functionality visible in the UI. However, applications need to perform the same access control checks on the server when each function is accessed. If requests are not verified, attackers will be able to forge requests in order to access functionality without proper authorization.
- "**A8—Cross-Site Request Forgery (CSRF):** A CSRF attack forces a logged-on victim's browser to send a forged HTTP request, including the victim's session cookie and any other automatically included authentication information, to a vulnerable web application. This allows the attacker to force the victim's browser to generate requests the vulnerable application thinks are legitimate requests from the victim.
- "**A9—Using Components with Known Vulnerabilities:** Components, such as libraries, frameworks, and other software modules, almost always run with full privileges. If a vulnerable component is exploited, such an attack can facilitate serious data loss or server takeover. Applications using components with known vulnerabilities may undermine application defences and enable a range of possible attacks and impacts.
- "**A10—Unvalidated Redirects and Forwards:** Web applications frequently redirect and forward users to other pages and websites, and use untrusted data to determine the destination pages. Without proper validation, attackers can redirect victims to phishing or malware sites, or use forwards to access unauthorized pages."[20]

CLOUD SECURE DATA LIFECYCLE

Data is the single most valuable asset for most organizations. Depending on the value of the information to their operations, security controls should be applied accordingly.

As with systems and other organizational assets, data should have a defined and managed lifecycle across the following key stages (*Figure 1.5*). According to Securosis, the data lifecycle is comprised of six phases, from creation to destruction.

- **Create:** New digital content is generated or existing content is modified.
- **Store:** Data is committed to a storage repository, which typically occurs directly after creation.
- **Use:** Data is viewed, processed, or otherwise used in some sort of activity (not including modification).
- **Share:** Information is made accessible to others—users, partners, customers, and so on.
- **Archive:** Data leaves active use and enters long-term storage.
- **Destroy:** Data is permanently destroyed using physical or digital means.

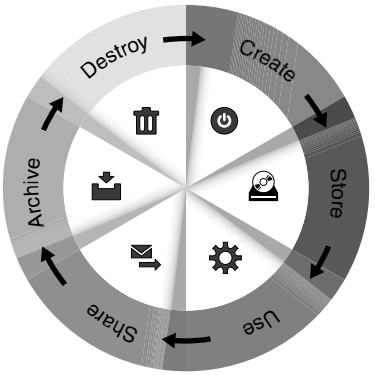

FIGURE 1.5 **Key stages of the data lifecycle.**

The lifecycle is not a single linear operation but a series of smaller lifecycles running in different environments. At all times, it is important to be aware of the logical and physical location of the data to satisfy audit, compliance, and other control requirements.

In addition to the location of the data, it is important to know who is accessing the data and how they are accessing it.

NOTE Different devices have specific security characteristics or limitations (BYOD, and so on).

INFORMATION AND DATA GOVERNANCE TYPES

Table 1.2 lists a sample of information and data governance types. Note that this may vary depending on your organization, geographic location, risk appetite, and so on.

TABLE 1.2 **Information and Data Governance Types**

FEATURE	DESCRIPTION
Information classification	What is the high-level description of valuable information categories (such as highly confidential, regulated)?
Information management policies	What activities are allowed for different information types?
Location and jurisdictional policies	Where can data be geographically located? What are the legal and regulatory implications or ramifications?
Authorizations	Who is allowed to access different types of information?
Custodianship	Who is responsible for managing the information at the bequest of the owner?

BUSINESS CONTINUITY AND DISASTER RECOVERY PLANNING

Business continuity management is the process by which risks and threats to the ongoing availability of services, business functions, and the organization are actively reviewed and managed at set intervals as part of the overall risk-management process. The goal is to keep the business operating and functioning in the event of a disruption.

Disaster recovery planning (DRP) is the process by which suitable plans and measures are taken to ensure that, in the event of a disaster (flood, storm, tornado, and so on), the business can respond appropriately with the view to recovering critical and essential operations (even somewhat limited) to a state of partial or full level of service in as little time as possible. The goal is to quickly establish, reestablish, or recover affected areas or elements of the business following a disaster.

Note that DR and business continuity are often confused or used interchangeably in some organizations. Wherever possible, be sure to use the correct terminology and highlight the differences between them.

Business Continuity Elements

From the perspective of the cloud customer, business continuity elements include the relevant security pillars of availability, integrity, and confidentiality.

The availability of the relevant resources and services is often the key requirement, along with the uptime and ability to access these on demand. Failure to ensure this results in significant impacts, including loss of earnings, loss of opportunities, and loss of confidence for the customer and provider.

Many security professionals struggle to keep their business continuity processes current once they have started to utilize cloud-based services. Equally, many fail to adequately update, amend, and keep their business continuity plans up to date in terms of complete coverage of services. This may be due to a number of factors; however, the key component contributing to this is that business continuity is operated mainly at set intervals and is not integrated fully into ongoing business operations. That is, business continuity activities are performed only annually or biannually, which may not take into account notable changes in business operations (such as the cloud) within relevant business units, sections, or systems.

Note that not all assets or services are equal! What are the key or fundamental components required to ensure the business or service can continue to be delivered? The answer to this question should shape and structure your business continuity and disaster recovery (BCDR) practices.

Critical Success Factors

Two critical success factors for business continuity when utilizing cloud-based services are as follows:

- Understand your responsibilities versus the CSP's responsibilities.
  - Customer responsibilities
  - CSP responsibilities
  - Understanding any interdependencies or third parties (supply chain risks)
  - Order of restoration (priority)
  - Appropriate frameworks and certifications held by the facility, services, and processes
  - Right to audit and make regular assessments of continuity capabilities
  - Communications of any issues or limited services
  - Identification of need for backups to be held onsite or offsite or with another CSP
- Clearly state and ensure the SLA addresses which components of business continuity and disaster recovery are covered and to what degree they are covered.
  - Penalties and compensation for loss of service
  - RTOs and RPOs
  - Loss of integrity or confidentiality
  - Points of contact and escalation processes

- Failover to maintain compliance
- Changes being communicated in a timely manner
- Clearly defined responsibilities
- Where usage of third parties is required per the agreed-upon SLA

The cloud customer should be in agreement with and fully satisfied with all the details relating to BCDR (including recovery times, responsibilities, and more) prior to signing any documentation or agreements that signify acceptance of the terms for system operation.

The customer typically pays for the associated time and costs of requesting amendments or changes to the relevant SLA.

Important SLA Components

Finally, regarding DR, the cloud customer should take a similar approach to ensure the following are fully understood and acted upon, prior to signing relevant SLAs and contracts:

- Undocumented single points of failure should not exist.
- Migration to alternate providers should be possible within agreed-upon timeframes.
- All components need to be supported by alternate CSPs in the event of a failover; if not, onsite services may be required as a fallback solution.
- Automated controls should be enabled to allow customers to verify data integrity.
- Where data backups are included, incremental backups should allow the user to select the desired settings, including coverage, frequency, and ease of use for recovery point restoration options.
- Regular assessment of the SLA and any changes that may affect the customer's ability to utilize cloud computing components for DR should be captured at regular and set intervals.

Although it's impossible to plan for every event or disaster that may occur, relevant plans and continuity measure should cover a number of logical groupings, which could be applied for unforeseen or unplanned incidents.

As cloud adoption and migration continue to expand, all affected or associated areas of business (technology and otherwise) should be reviewed under BCDR plans, thus ensuring that any changes for the customer or provider are captured and acted upon. Imagine the challenges of trying to restore or act upon a loss of availability, when processes, controls, or technologies have changed without the plans being updated or amended to reflect such changes.

The following ISO/IEC documents may be of use to CCSPs as they are considering what items an SLA will need to address:

- ISO/IEC DIS 19086-1, "Information Technology—Cloud Computing—Service Level Agreement (SLA) Framework—Part 1: Overview and Concepts"
- ISO/IEC NP 19086-2, "Information Technology—Cloud Computing—Service Level Agreement (SLA) Framework and Technology—Part 2: Metrics"
- ISO/IEC CD 19086-3, "Information Technology –Cloud Computing—Service Level Agreement (SLA) Framework and Technology—Part 3: Core Requirements"
- ISO/IEC AWI 19941, "Information Technology –Cloud Computing—Interoperability and Portability"
- ISO/IEC CD 19944, "Information Technology—Cloud Computing—Data and Their Flow Across Devices and Cloud Services"
- ISO/IEC FDIS 20933, "Information Technology—Distributed Application Platforms and Services (DAPS)—Access Systems"

COST-BENEFIT ANALYSIS

Cost is often identified as a key driver for the adoption of cloud computing. The challenge with decisions being made solely or exclusively on cost savings can come back to haunt the organization or entity that failed to take a risk-based view and factor in the relevant effects that may materialize.

- **Resource pooling:** Resource sharing is essential to the attainment of significant cost savings when adopting a cloud computing strategy. This is often coupled with pooled resources being used by different consumer groups at different times.
- **Shift from CapEx to OpEx:** The shift from capital expenditure (CapEx) to operational expenditure (OpEx) is seen as a key factor for many organizations as their requirement to make significant purchases of systems and resources is minimized. Given the constant evolution of technology and computing power, memory, capabilities, and functionality, many traditional systems purchased lose value almost instantly.
- **Factor in time and efficiencies:** Given that organizations rarely acquire used technology or servers, almost all purchases are of new and recently developed technology. But it's not just technology investment savings. Time and efficiencies achieved can be the greatest savings achieved when utilizing cloud computing.
- **Include depreciation:** When you purchase a new car, the value deteriorates the moment you drive the car off the showroom floor. The same applies for IT, only

with newer and more desirable cars, technologies, and models being released every few months or years. Using this analogy clearly highlights why so many organizations are now opting to lease cloud services as opposed to constantly investing in technologies that become outdated in relatively short periods.

- **Reduction in maintenance and configuration time:** Remember all those days, weeks, months, and years spent maintaining, operating, patching, updating, supporting, engineering, rebuilding, and generally making sure everything needed was done to the systems and applications required by the business users? Well, given that the CSP now handles a large portion of those duties (if not all—depending on which cloud service you are using), the ability to free up, utilize, and reallocate resources to other technology or related tasks could prove to be invaluable.
- **Shift in focus:** Technology and business personnel being able to focus on the key elements of their role, instead of the daily firefighting and responding to issues and technology components, comes as a welcome change to those professionals serious about their functions.
- **Utilities costs:** Outside of the technology and operational elements, from a utilities cost perspective, massive savings can be achieved with the reduced requirement for power, cooling, support agreements, data center space, racks, cabinets, and so on. Large organizations that have migrated big portions of the data center components to cloud-based environments have reported tens of thousands to hundreds of thousands in direct savings from the utilities elements. Green IT is very much at the fore of many global organizations, and cloud computing plays toward that focus in a strong way.
- **Software and licensing costs:** Software and relevant licensing costs present a major cost saving as well because you only pay for the licensing used versus the bulk or enterprise licensing levels of traditional non-cloud-based infrastructure models.
- **Pay per usage:** As outlined by the CapEx versus OpEx discussion earlier in this section, cloud computing gives businesses a new and clear benefit: pay per usage. In terms of traditional IT functions, when systems and infrastructure assets were acquired, they were seen as a "necessary or required spend" for the organization; however, with cloud computing, they can now be monitored, categorized, and billed to specified functions or departments based on usage. This is a significant win and driver for IT departments because it releases pressure to reduce spending and allows for billing of usage for relevant cost bases directly to those, as opposed to absorbing the costs themselves as a business requirement.

With departments and business units now able to track costs and usage, it's easy to work out the amount of money spent versus the amount saved in traditional type computing. Sounds pretty straightforward, right?

- **Other factors:** What about new technologies, new or revised roles, legal costs, contract and SLA negotiations, additional governance requirements, training required, CSP interactions, and reporting? All these may impact and alter the price you see versus the price you pay, otherwise known as the total cost of ownership (TCO).

Many organizations have not factored in such costs to date. As such, their view of cost savings may be skewed or misguided somewhat.

CERTIFICATION AGAINST CRITERIA

If it cannot be measured, it cannot be managed.

This is a statement that any auditor and security professional should abide by regardless of his focus. How can someone have confidence, awareness, and assurances that he and the CSP are taking the correct steps to ensure that data is secured properly? Frameworks and standards hold the key here.

Why are users and entities still unconvinced that cloud computing is a good option, particularly from a security perspective? The reason is simple: no international cloud computing standards or security standards exist.

In the absence of cloud-specific security standards that are universally accepted by providers and customers alike, you'll deal with a patchwork of security standards, frameworks, and controls that are being applied to cloud environments. These include but are not limited to the following:

- ISO/IEC 27001:2013[21]
- ISO/IEC 27002:2013
- ISO/IEC 27017:2015
- SOC 1/SOC 2/SOC 3
- NIST SP 800-53
- PCI DSS

Possibly the most widely known and accepted information security standard, ISO 27001 was originally developed and created by the British Standards Institute, under the name of BS 7799. The standard was adopted by the International Organization for Standardization (ISO) and rebranded ISO 27001. ISO 27001 is the standard to which

organizations certify, as opposed to ISO 27002, which is the best practice framework to which many others align.

ISO 27001:2005 consisted of 133 controls across 11 domains of security, focusing on the protection of information assets in their various forms (digital, paper, and so on). Since September 2013, ISO 27001 has been updated to ISO 27001:2013 and now consists of 35 control objectives and 114 controls spread over 14 domains.

Domains include these:

1. Information Security Policies
2. Organization of Information Security
3. Human Resources Security
4. Asset Management
5. Access Control
6. Cryptographic
7. Physical and Environmental Security
8. Operations Security
9. Communications Security
10. System Acquisition, Development, and Maintenance
11. Supplier Relationship
12. Information Security Incident Management
13. Information Security Business Continuity Management
14. Compliance

By its nature, ISO 27001 is designed to be vendor and technology agnostic (that is, it does not view them differently). As such, it looks for the information security management system (ISMS) to address the relevant risks and components in a manner that is appropriate and adequate based on the risks.

Even though ISO 27001 is the most advanced security standard widely used today, it does not specifically look at the risks associated with cloud computing. As such, it cannot be deemed as fully comprehensive when measuring security in cloud-based environments.

All standards and frameworks assist in the structure and standardization of security practices; however, they cannot be applied across multiple environments (of differing natures), deployments, and other components with 100 percent confidence and completeness, given the variations and specialized elements associated with cloud computing.

Due to its importance overall, ISO 27001 will continue to be used by CSPs and required by cloud customers as one of the key security frameworks for cloud environments.

ISO/IEC 27002:2013

ISO/IEC 27002:2013 provides guidelines for organizational information security standards including the selection, implementation, and management of controls taking into consideration the organization's information security risk environments. It is designed to be used by organizations that intend to select controls within the process of implementing an ISMS based on ISO/IEC 27001. It can also be used by organizations to implement commonly accepted information security controls and develop their own information security management guidelines.

ISO/IEC 27017:2015

ISO/IEC 27017:2015 offers guidelines for information security controls applicable to the provision and use of cloud services by providing additional implementation guidance for relevant controls specified in ISO/IEC 27002 and additional controls with implementation guidance that specifically relate to cloud services. ISO 27017 provides controls and implementation guidance for both CSPs and cloud service customers.

SOC 1/SOC 2/SOC 3[22]

The Statement on Auditing Standards 70 (SAS 70) was replaced by Service Organization Control (SOC) Type 1 and Type 2 reports in 2011 following changes and a more comprehensive approach to auditing being demanded by customers and clients alike. For years, SAS 70 was seen as the de facto standard for data center customers to obtain independent assurance that their data center service provider had effective internal controls in place for managing the design, implementation, and execution of customer information.

SAS 70 consisted of Type 1 and Type 2 audits. The Type 1 audit was designed to assess the sufficiency of the service provider's controls as of a particular date, and the Type 2 audit was designed to assess the effectiveness of the controls as of a certain date (point-in-time assessment).

Like many other frameworks, SAS 70 audits focused on verifying that the controls had been implemented and followed but did not focus on the overall completeness or effectiveness of the controls implemented. Think of having an alarm but not checking whether it was effective, functioning, or correctly installed.

SOC reports are performed in accordance with Statement on Standards for Attestation Engagements (SSAE) 16, Reporting on Controls at a Service Organization.

- SOC 1 reports focus solely on controls at a service provider that are likely to be relevant to an audit of a subscriber's financial statements.
- SOC 2 and SOC 3 reports address controls of the service provider that relate to operations and compliance.

There are some key distinctions between SOC 1, SOC 2, and SOC 3:

SOC 1 SOC 1 reports can be one of two types:

- A Type 1 report presents the auditors' opinion regarding the accuracy and completeness of management's description of the system or service as well as the suitability of the design of controls as of a specific date.
- Type 2 reports include the Type 1 criteria and audit of the operating effectiveness of the controls throughout a declared period, generally between 6 months and 1 year.

SOC 2 SOC 2 reporting was specifically designed for IT-managed service providers and cloud computing. The report specifically addresses any number of the five so-called Trust Services principles, which follow:

- **Security:** The system is protected against unauthorized access, both physical and logical.
- **Availability:** The system is available for operation and use as committed or agreed.
- **Processing Integrity:** System processing is complete, accurate, timely, and authorized.
- **Confidentiality:** Information designated as confidential is protected as committed or agreed.
- **Privacy:** Personal information is collected, used, retained, disclosed, and disposed of in conformity with the provider's privacy policy.

SOC 3 Reporting also uses the Trust Services principles but provides only the auditor's report on whether the system achieved the specified principle, without disclosing relevant details and sensitive information.

A key difference between a SOC 2 report and a SOC 3 report is that a SOC 2 report is generally restricted in distribution and coverage (due to the information it contains), and a SOC 3 report is broadly available, with limited information and details included within it (often used to instill confidence in perspective clients or for marketing purposes).

To review:

- **SOC 1:** This is intended for those interested in financial statements.
- **SOC 2:** Information technology personnel will be interested.
- **SOC 3:** This is used to illustrate conformity, compliance, and security efforts to current or potential subscribers and customers of cloud services.

NIST SP 800-53

NIST is an agency of the U.S. government that makes measurements and sets standards as needed for industry or government programs. The primary goal and objective of the 800-53[23] standard is to ensure that appropriate security requirements and security controls are applied to all U.S. federal government information and information management systems.

The standard requires that risk be assessed and the determination made whether additional controls are needed to protect organizational operations (including mission, functions, image, or reputation), organizational assets, individuals, other organizations, or the nation.

The 800-53 standard—"Security and Privacy Controls for Federal Information Systems and Organizations"—underwent its fourth revision in April 2013.

Primary updates and amendments include these:

- Assumptions relating to security control baseline development
- Expanded, updated, and streamlined tailoring guidance
- Additional assignment and selection statement options for security and privacy controls
- Descriptive names for security and privacy control enhancements
- Consolidated security controls and control enhancements by family with baseline allocations
- Tables for security controls that support development, evaluation, and operational assurance
- Mapping tables for international security standard ISO/IEC 15408 (Common Criteria)

Although the NIST Risk Management Framework provides the pieces and parts for an effective security program, it is aimed at government agencies focusing on the following key components:

2.1 Multitiered Risk Management

2.2 Security Control Structure

2.3 Security Control Baselines

2.4 Security Control Designations

2.5 External Service Partners

2.6 Assurance and Trustworthiness

2.7 Revisions and Extensions

3.1 Selecting Security Control Baselines

3.2 Tailoring Security Control Baselines

3.3 Creating Overlays

3.4 Document the Control Selection Process
3.5 New Development and Legacy Systems

One major issue that corporate security teams encounter when trying to base a program on the NIST SP 800-53 Risk Management Framework is that publicly traded organizations are not bound by the same security assumptions and requirements as government agencies. Government organizations are established to fulfill legislated missions and are required to collect, store, manipulate, and report sensitive data. Finally, a large percentage of these activities in a publicly traded organization is governed by cost-benefit analysis, boards of directors, and shareholder opinion, as opposed to government direction and influence.

For those looking to understand the similarities and overlaps with NIST SP 800-53 and ISO 27001/2, there is a mapping matrix listed within the 800-53 Revision 4 document.

PCI DSS

Visa, MasterCard, and American Express established PCI DSS[24] as a security standard to which all organizations or merchants that accept, transmit, or store cardholder data, regardless of size or number of transactions, must comply.

PCI DSS was established following a number of significant credit card breaches. It is a comprehensive and intensive security standard that lists both technical and nontechnical requirements based on the number of credit card transactions for the applicable entities.

Merchant Levels Based on Transactions

Table 1.3 illustrates the various merchant levels based on the number of transactions.

TABLE 1.3 Merchant Levels Based on Transactions

MERCHANT LEVEL	DESCRIPTION
1	Any merchant—regardless of acceptance channel—processing more than 6 million transactions per year. Any merchant that the credit card issuer, at its sole discretion, determines should meet the Level 1 merchant requirements to minimize risk to the credit card issuer's system.
2	Any merchant—regardless of acceptance channel—processing 1–6 million credit card transactions per year.
3	Any merchant processing 20,000 to 1 million credit card e-commerce transactions per year.
4	Any merchant processing fewer than 20,000 credit card e-commerce transactions per year and all other merchants—regardless of acceptance channel—processing up to 1 million credit card transactions per year.

For specific information and requirements, be sure to check with the PCI Security Standard Council.

Merchant Requirements

All merchants, regardless of level and relevant service providers, are required to comply with the following 12 domains/requirements:

- Install and maintain a firewall configuration to protect cardholder data.
- Avoid using vendor-supplied defaults for system passwords and other security parameters.
- Protect stored cardholder data.
- Encrypt transmission of cardholder data across open, public networks.
- Use and regularly update antivirus software.
- Develop and maintain secure systems and applications.
- Restrict access to cardholder data by business need-to-know.
- Assign a unique ID to each person with computer access.
- Restrict physical access to cardholder data.
- Track and monitor all access to network resources and cardholder data.
- Regularly test security systems and processes.
- Maintain a policy that addresses information security.

The 12 requirements list more than 200 controls that specify required and minimum security requirements for the merchants and service providers to meet their compliance obligations.

Failure to meet and satisfy the PCI DSS requirements (based on merchant level and processing levels) can result in significant financial penalties, suspension of credit cards as a payment channel, escalation to a higher merchant level, and potentially greater assurances and compliance requirements in the event of a breach in which credit card details may have be compromised or disclosed.

Since its establishment, PCI DSS has undergone a number of significant updates, through to the current version.

Due to the more technical and more black-and-white nature of its controls, many see PCI DSS as a reasonable and sufficient technical security standard. People believe that if it is good enough to protect their credit card and financial information, it should be a good baseline for cloud security.

SYSTEM AND SUBSYSTEM PRODUCT CERTIFICATION

System and subsystem product certification is used to evaluate the security claims made for a system and its components. Although there have been several evaluation frameworks available for use over the years, such as the Trusted Computer System Evaluation Criteria (TCSEC) developed by the United States Department of Defense, the Common Criteria (CC), discussed next, is the one that is internationally accepted and used most often.

CC

The CC[25] is an international set of guidelines and specifications (ISO/IEC 15408) developed for evaluating information security products, with the view to ensuring they meet an agreed-upon security standard for government entities and agencies.

CC Components

Officially, the CC is known as the "Common Criteria for Information Technology Security Evaluation." Until 2005, it was known as "The Trusted Computer System Evaluation Criteria." The CC is updated periodically.

Distinctly, the CC has two key components:

- **Protection profiles:** Define a standard set of security requirements for a specific type of product, such as a firewall, IDS, or unified threat management (UTM).
- **The evaluation assurance levels (EALs):** Define how thoroughly the product is tested. EALs are rated using a sliding scale from 1–7, with 1 being the lowest-level evaluation and 7 being the highest.

  The higher the level of evaluation, the more quality assurance (QA) tests the product would have undergone.

NOTE Undergoing more tests does not necessarily mean the product is more secure.

The seven EALs are as follows:

- **EAL1:** Functionally tested
- **EAL2:** Structurally tested
- **EAL3:** Methodically tested and checked
- **EAL4:** Methodically designed, tested, and reviewed
- **EAL5:** Semiformally designed and tested
- **EAL6:** Semiformally verified design and tested
- **EAL7:** Formally verified design and tested

CC Evaluation Process

The goal of CC certification is to ensure customers that the products they are buying have been evaluated and that a vendor-neutral third party has verified the vendor's claims.

To submit a product for evaluation, follow these steps:

1. The vendor must complete a Security Target (ST) description that provides an overview of the product's security features.
2. A certified laboratory then tests the product to evaluate how well it meets the specifications defined in the protection profile.
3. A successful evaluation leads to an official certification of the product.

Note that CC looks at certifying a product only and does not include administrative or business processes.

FIPS 140-2

To maintain ongoing confidentiality and integrity of relevant information and data, you can use encryption and cryptography as a primary choice, specifically in various cloud computing deployment service types.

Federal Information Processing Standard (FIPS)[26] 140 Publication Series was issued by NIST to coordinate the requirements and standards for cryptography modules covering both hardware and software components for cloud and traditional computing environments.

The FIPS 140-2 standard provides four distinct levels of security intended to cover a range of potential applications and environments with emphasis on secure design and implementation of a cryptographic module.

Relevant specifications include these:

- Cryptographic module specification
- Cryptographic module ports
- Interfaces, roles, and services
- Authentication
- Physical security
- Operational environment
- Cryptographic key management
- Design assurance
- Controls and mitigating techniques against attacks

FIPS 140-2 Goal

The primary goal for the FIPS 140-2 standard is to accredit and distinguish secure and well-architected cryptographic modules produced by private sector vendors who seek to or are in the process of having their solutions and services certified for use in U.S. government departments and regulated industries (this includes financial services and healthcare) that collect, store, transfer, or share data that is deemed to be sensitive but not classified (that is, top secret).

Finally, when assessing the level of controls, FIPS is measured using a Level 1 to Level 4 rating. Despite the ratings and their associated requirements, FIPS does not state what level of certification is required by specific systems, applications, or data types.

FIPS Levels

The breakdown of the levels follows:

- **Security Level 1:** The lowest level of security. To meet Level 1 requirements, basic cryptographic module requirements are specified for at least one approved security function or approved algorithm. Encryption of a PC board presents an example of a Level 1 rating.
- **Security Level 2:** Enhances the required physical security mechanisms listed within Level 1 and requires that capabilities exist to illustrate evidence of tampering, including locks that are tamper proof on perimeter and internal covers to prevent unauthorized physical access to encryption keys.
- **Security Level 3:** Looks to develop the basis of Level 1 and Level 2 to include preventing the intruder from gaining access to information and data held within the cryptographic module. Additionally, physical security controls required at Level 3 should move toward detecting access attempts and responding appropriately to protect the cryptographic module.
- **Security Level 4:** Represents the highest rating. Security Level 4 provides the highest level of security, with mechanisms providing complete protection around the cryptographic module with the intent of detecting and responding to all unauthorized attempts at physical access. Upon detection, immediate zeroization of all plaintext critical security parameters (*also known as CSPs but not to be confused with cloud service providers*).[27] Security Level 4 undergoes rigid testing to ensure its adequacy, completeness, and effectiveness.

All testing is performed by accredited third-party laboratories and is subject to strict guidelines and quality standards. Upon completion of testing, all ratings are provided, along with an overall rating on the vendor's independent validation certificate.

From a cloud computing perspective, these requirements form a necessary and required baseline for all U.S. government agencies that may be looking to utilize or avail cloud-based services. Outside of the United States, FIPS does not typically act as a driver or a requirement; however, other governments and enterprises tend to recognize the FIPS validation as an enabler or differentiator over other technologies that have not undergone independent assessments or certification.

SUMMARY

Cloud computing covers a wide range of topics focused on the concepts, principles, structures, and standards used to monitor and secure assets and those controls used to enforce various levels of AIC across IT services throughout the enterprise. Security practitioners focused on cloud security must use and apply standards to ensure that the systems under their protection are maintained and supported properly. Today's environment of highly interconnected, interdependent systems necessitates the requirement to understand the linkage between information technology and meeting business objectives. Information security management communicates the risks accepted by the organization due to the currently implemented security controls and continually works to cost effectively enhance the controls to minimize the risk to the company's information assets.

REVIEW QUESTIONS

1. Which of the following are attributes of cloud computing?
   - **A.** Minimal management effort and shared resources
   - **B.** High cost and unique resources
   - **C.** Rapid provisioning and slow release of resources
   - **D.** Limited access and service provider interaction
2. Which of the following are distinguishing characteristics of a managed service provider?
   - **A.** Have some form of a NOC but no help desk.
   - **B.** Be able to remotely monitor and manage objects for the customer and reactively maintain these objects under management.
   - **C.** Have some form of a help desk but no NOC.
   - **D.** Be able to remotely monitor and manage objects for the customer and proactively maintain these objects under management.
3. Which of the following are cloud computing roles?
   - **A.** Cloud customer and financial auditor
   - **B.** CSP and backup service provider
   - **C.** Cloud service broker and user
   - **D.** Cloud service auditor and object
4. Which of the following are essential characteristics of cloud computing? (Choose two.)
   - **A.** On-demand self-service
   - **B.** Unmeasured service
   - **C.** Resource isolation
   - **D.** Broad network access
5. Which of the following are considered to be the building blocks of cloud computing?
   - **A.** Data, access control, virtualization, and services
   - **B.** Storage, networking, printing, and virtualization
   - **C.** CPU, RAM, storage, and networking
   - **D.** Data, CPU, RAM, and access control

6. When using an IaaS solution, what is the capability provided to the customer?
   - **A.** To provision processing, storage, networks, and other fundamental computing resources where the consumer is not able to deploy and run arbitrary software, which can include OSs and applications
   - **B.** To provision processing, storage, networks, and other fundamental computing resources where the provider is able to deploy and run arbitrary software, which can include OSs and applications
   - **C.** To provision processing, storage, networks, and other fundamental computing resources where the auditor is able to deploy and run arbitrary software, which can include OSs and applications
   - **D.** To provision processing, storage, networks, and other fundamental computing resources where the consumer is able to deploy and run arbitrary software, which can include OSs and applications
7. When using an IaaS solution, what is a key benefit provided to the customer?
   - **A.** Metered and priced usage on the basis of units consumed
   - **B.** The ability to scale up infrastructure services based on projected usage
   - **C.** Increased energy and cooling system efficiencies
   - **D.** Transferred cost of ownership
8. When using a PaaS solution, what is the capability provided to the customer?
   - **A.** To deploy onto the cloud infrastructure provider-created or acquired applications created using programming languages, libraries, services, and tools that the provider supports. The consumer does not manage or control the underlying cloud infrastructure, including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.
   - **B.** To deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools that the provider supports. The provider does not manage or control the underlying cloud infrastructure, including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.
   - **C.** To deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools that the provider supports. The consumer does not manage or control the underlying cloud infrastructure, including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.

D. To deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools that the consumer supports. The consumer does not manage or control the underlying cloud infrastructure, including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.

9. What is a key capability or characteristic of PaaS?

   A. Support for a homogenous hosting environment

   B. Ability to reduce lock-in

   C. Support for a single programming language

   D. Ability to manually scale

10. When using a SaaS solution, what is the capability provided to the customer?

    A. 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 (for example, web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure, including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

    B. 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 (for example, web-based email), or a program interface. The consumer does manage or control the underlying cloud infrastructure, including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

    C. To use the consumer'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 (for example, web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure, including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

    D. To use the consumer'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 (for example, web-based email), or a program interface. The consumer does manage or control the underlying cloud infrastruc-

ture, including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

11. What are the four cloud deployment models?
    - **A.** Public, internal, hybrid, and community
    - **B.** External, private, hybrid, and community
    - **C.** Public, private, joint, and community
    - **D.** Public, private, hybrid, and community
12. What are the six stages of the cloud secure data lifecycle?
    - **A.** Create, use, store, share, archive, and destroy
    - **B.** Create, store, use, share, archive, and destroy
    - **C.** Create, share, store, archive, use, and destroy
    - **D.** Create, archive, use, share, store, and destroy
13. What are SOC 1/SOC 2/SOC 3?
    - **A.** Risk management frameworks
    - **B.** Access controls
    - **C.** Audit reports
    - **D.** Software development phases
14. What are the five Trust Services principles?
    - **A.** Security, Availability, Processing Integrity, Confidentiality, and Privacy
    - **B.** Security, Auditability, Processing Integrity, Confidentiality, and Privacy
    - **C.** Security, Availability, Customer Integrity, Confidentiality, and Privacy
    - **D.** Security, Availability, Processing Integrity, Confidentiality, and Nonrepudiation
15. What is a security-related concern for a PaaS solution?
    - **A.** Virtual machine attacks
    - **B.** Web application security
    - **C.** Data access and policies
    - **D.** System and resource isolation

NOTES

[1] http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf (p. 6)

[2] http://www.mspalliance.com/

[3] *Governance Reimagined: Organizational Design, Risk and Value Creation*, by David R. Koenig, John Wiley & Sons, Inc., p. 160.

[4] http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf (p. 7)

[5] http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf (p. 6)

[6] http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf (p. 6)

[7] http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf (p. 7)

[8] http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf (p. 7)

[9] http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf (p. 7)

[10] http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf (p. 7)

[11] http://www.sabsa.org/

[12] https://www.axelos.com/itil

[13] http://www.opengroup.org/subjectareas/enterprise/togaf

[14] http://www.opengroup.org/subjectareas/platform3.0/cloudcomputing

[15] See the following for the October 22, 2014 announcement by NIST of the final publication release of the roadmap: http://www.nist.gov/itl/antd/cloud-102214.cfm

[16] See the following for the LDAP X.500 RFC: https://tools.ietf.org/html/rfc2247

[17] https://downloads.cloudsecurityalliance.org/initiatives/top_threats/The_Notorious_Nine_Cloud_Computing_Top_Threats_in_2013.pdf

[18] http://www.cert.org/insider-threat/

[19] See the following for more information: https://cloudsecurityalliance.org/research/ccm/

[20] https://www.owasp.org/index.php/Category:OWASP_Top_Ten_Project

[21] http://www.iso.org/iso/catalogue_detail?csnumber=54534

[22] https://www.ssae-16.com/

[23] http://nvlpubs.nist.gov/nistpubs/SpecialPublications/NIST.SP.800-53r4.pdf

[24] https://www.pcisecuritystandards.org/documents/PCI_DSS_v3.pdf

[25] http://www.commoncriteriaportal.org/files/ccfiles/CCPART1V3.1R4.pdf

[26] http://csrc.nist.gov/groups/STM/cmvp/standards.html

[27] In cryptography, zeroization is the practice of erasing sensitive parameters (electronically stored data, cryptographic keys, and CSPs) from a cryptographic module to prevent their disclosure if the equipment is captured.