Chapter 1

Architectural Concepts

THE OBJECTIVE OF THIS CHAPTER IS TO ACQUAINT THE READER WITH THE FOLLOWING CONCEPTS:

✓ **Domain 1: Cloud Concepts, Architecture, and Design**

- 1.1. Understand Cloud Computing Concepts
  - 1.1.1. Cloud Computing Definitions
  - 1.1.2. Cloud Computing Roles and Responsibilities (e.g., cloud service customer, cloud service provider, cloud service partner, cloud service broker, regulator)
  - 1.1.3. Key Cloud Computing Characteristics (e.g., on-demand self-service, broad network access, multitenancy, rapid elasticity and scalability, resource pooling, measured service)
  - 1.1.4. Building Block Technologies (e.g., virtualization, storage, networking, databases, orchestration)
- 1.2. Describe Cloud Reference Architecture
  - 1.2.1. Cloud Computing Activities
  - 1.2.2. Cloud Service Capabilities (e.g., application capability types, infrastructure capability types)
  - 1.2.3. Cloud Service Categories (e.g., Software as a Service (SaaS), Infrastructure as a Service (IaaS), Platform as a Service (PaaS))
  - 1.2.4. Cloud Deployment Models (e.g., public, private, hybrid, community, multi-cloud)
  - 1.2.5. Cloud Shared Considerations (e.g., interoperability, portability, reversibility, availability, security, privacy, resiliency, performance, governance, maintenance and versioning, service levels and service-level agreements (SLA), auditability, regulatory, outsourcing)

- 1.2.6. Impact of Related Technologies (e.g., data science, machine learning, artificial intelligence (AI), blockchain, Internet of Things (IoT), containers, quantum computing, edge computing, confidential computing, DevSecOps)

- 1.3. Understand Design Principles of Secure Cloud Computing
  - 1.4.3. Business Impact Analysis (BIA) (e.g., cost-benefit analysis, return on investment (ROI))

Cloud computing is everywhere. The modern business depends upon a wide variety of software, platforms, and infrastructure hosted in the cloud, and security professionals must understand how to protect the information and resources used by their organizations, wherever those assets reside.

In this chapter, we introduce the basic concepts of cloud computing and help you understand the foundational material you'll need to know as you begin your journey toward the Certified Cloud Security Professional (CCSP) certification.

Cloud Characteristics

Cloud computing is the most transformative development in information technology in the past decade. Organizations around the world are retooling their entire IT strategies to embrace the cloud, and this change is causing disruptive impact across all sectors of technology.

But what is the cloud? Let's start with a simple definition: cloud computing is any case where a provider is delivering computing to a customer at a remote location over a network. This definition is broad and encompasses many different types of activity.

There are some common characteristics that we use to define cloud computing:

- Broad network access
- On-demand self-service
- Resource pooling
- Rapid elasticity and scalability
- Measured, or "metered," service

These traits are expressed succinctly in the NIST definition of cloud computing.

NIST 800-145 Cloud Computing Definition

The official NIST definition of cloud computing says, "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."

These characteristics are also similar to how cloud computing is defined in ISO 17788 (`www.iso.org/iso/catalogue_detail?csnumber=60544`).

Let's explore these characteristics in more detail.

- **Broad network access** means services are consistently accessible over the network. We might access them by using a web browser or Secure Shell (SSH) connection, but the general idea is that no matter where we or our users are physically located, we can access resources in the cloud.
- **On-demand self-service** refers to the model that allows customers to scale their compute and/or storage needs with little or no intervention from or prior communication with the provider. This means that technologists can access cloud resources almost immediately when they need them to do their jobs. That's an incredible increase in agility for individual contributors and, by extension, the organization. Before the era of on-demand computing, a technologist who wanted to try out a new idea might have to spec out the servers required to implement the idea, gain funding approval, order the hardware, wait for it to arrive, physically install it, and configure an operating system before getting down to work. That might have taken weeks, while today, the same tasks can be accomplished in the cloud in a matter of seconds. On-demand self-service computing is a true game changer.
- **Resource pooling** is the characteristic that allows the cloud provider to meet various demands from customers while remaining financially viable. The cloud provider can make capital investments that greatly exceed what any single customer could provide on their own and can apportion these resources as needed so that the resources are not underutilized (which would mean a wasteful investment) or overtaxed (which would mean a decrease in level of service).
- **Rapid elasticity and scalability** allows the customer to grow or shrink the IT footprint (number of users, number of machines, size of storage, and so on) as necessary to meet operational needs without excess capacity. In the cloud, this can be done in moments as opposed to the traditional environment, where acquisition and deployment of resources (or dispensing old resources) can take weeks or months. In many cases, this scaling can occur automatically, using code to add and remove resources as demands change.
- **Measured service**, or metered service, means that almost everything you do in the cloud is metered. Cloud providers measure the number of seconds you use a virtual server, the amount of disk space you consume, the number of function calls you make, and many other measures. This allows them to charge you for precisely the services you use—no more and no less. This is the same model commonly used by public utilities providing commodity services such as electricity and water. The measured service model is a little intimidating when you first encounter it, but it provides cloud customers with the ability to manage their utilization effectively and achieve the economic benefits of the cloud.

Real World Scenario

Online Shopping

Think of retail demand during the pre-holiday rush toward the end of the year. The sheer volume of customers and transactions greatly exceeds all normal operations throughout the rest of the year. When this happens, retailers who offer online shopping can see great benefit from hosting their sales capability in the cloud. The cloud provider can apportion resources necessary to meet this increased demand and will charge for the increased usage at a negotiated rate, but when shopping drops off after the holiday, the retailers will not continue to be charged at the higher rate.

Elasticity vs. Scalability

Many people use the terms *elasticity* and *scalability* interchangeably, but they are actually subtly different concepts.

Strictly speaking, *scalability* refers to the ability of a system to grow as demand increases. This growth does not need to be automated, but it does need to be possible. Scalability may come from using the automated scaling features of a cloud provider, or it may come from adding physical hardware to a system.

Elasticity refers to the ability of a system to dynamically grow and shrink based upon the current level of demand. Administrators may set up a system to automatically add storage, processing power, or network capacity as demand increases and then release those resources when demand is lower. This provides tremendous cost efficiency by only purchasing expensive computing resources when they are actually needed.

Business Requirements

In most businesses, the IT department is not a profit center; it provides a support function that allows other business units to generate a profit. Cybersecurity teams definitely fit into this category—they generally don't do anything that generates revenue for the business, and from the perspective of business leaders, they represent a sunk cost that reduces efficiency

by lowering profits. In fact, security activities often hinder business efficiency (because, generally, the more secure something is, be it a device or a process, the less efficient it will be). This is why the business needs of the organization drive security decisions and not the other way around.

A successful organization will gather as much information about operational business requirements as possible; this information can be used for many purposes, including several functions in the security realm. (We'll touch on this throughout the book, but a few examples include the business continuity and disaster recovery effort, the risk management plan, and data categorization.) Likewise, the astute security professional needs to understand as much as possible about the operation of the organization. Operational aspects of the organization can help security personnel better perform their tasks no matter what level or role they happen to be assigned to. Consider the following examples:

- A network security administrator has to know what type of traffic to expect based on the business of the organization.
- The intrusion detection analyst has to understand what the organization is doing, how business activities occur, and where (geographically) the business is operating to better understand the nature and intensity of potential external attacks and how to adjust baselines accordingly.
- The security architect has to understand the various needs of the organizational departments to enhance their operation without compromising their security profile.
- Security leaders must not only understand the technologies used by the organization but also the associated risks and how to appropriately manage them.

Functional requirements: Those performance aspects of a device, process, or employee that are necessary for the business task to be accomplished. Example: A salesperson in the field must be able to connect to the organization's network remotely.

Nonfunctional requirements: Those aspects of a device, process, or employee that are not necessary for accomplishing a business task but are desired or expected. Example: The salesperson's remote connection must be secure.

As organizations consider their distribution of resources between the cloud and on-premises computing environments, they must select a mix that is appropriate for their needs. This is not a decision made lightly, and the business requirements must be supported by this transition. There are also different cloud service and delivery models of cloud computing, and an organization must decide which one will optimize success.

Understanding the Existing State

A true evaluation and understanding of the business processes, assets, and requirements are essential. Failing to properly capture the full extent of the business needs could result in not having an asset or capability in the new environment after migration to the cloud.

At the start of this effort, however, the intent is not to determine what will best fulfill the business requirements but to determine what those requirements are. A full inventory of assets, processes, and requirements is necessary, and there are various methods for collecting this data. Typically, several methods are used jointly as a means to reduce the possibility of missing something.

Here are some possible methods for gathering business requirements:

- Interviewing functional managers
- Interviewing users
- Interviewing senior management
- Observing employees doing their jobs
- Surveying customers
- Collecting network traffic
- Inventorying assets
- Collecting financial records
- Collecting insurance records
- Collecting marketing data
- Collecting regulatory mandates

After sufficient data has been collected, a detailed analysis is necessary. This is the point where a *business impact analysis (BIA)* takes place.

The BIA is an assessment of the priorities given to each asset and process within the organization. A proper analysis should consider the effect (impact) any harm to or loss of each asset might mean to the organization overall. During the BIA, special care should be paid to identifying critical paths and single points of failure. You also need to determine the costs of compliance—that is, the legislative and contractual requirements mandated for your organization. Your organization's regulatory restrictions will be based on many variables, including the jurisdictions where your organization operates, the industry the organization is in, the types and locations of your customers, and so on.

Assets can be tangible or intangible. They can include hardware, software, intellectual property, personnel, processes, and so on. An example of tangible assets would be things like routers and servers, whereas intangible assets are generally something you cannot touch, such as software code, expressions of ideas, and business methodologies.

Cost/Benefit Analysis

Once you have a clear picture of what your organization does in terms of lines of business and processes, you can get a better understanding of what benefits the organization might derive from cloud migration as well as the costs associated with the move. Conducting a *cost/benefit analysis* helps you understand this trade-off in clear financial terms.

Obviously, the greatest driver pushing organizations toward cloud migration at the moment is perceived cost savings, and that is a significant and reasonable consideration. The next few sections describe some aspects of that consideration.

Reduction in Capital Expenditure

If your organization buys a device for use in its internal environment, the capacity of that device will either be fully utilized or (more likely) not. If the device is used at its fullest capacity, then it's quite likely that the function for which it is needed may experience inefficiencies at some point. Even a small uptick in demand for that device will overload its capacity. However, if the device is not fully utilized, then the organization has paid for something for which it is getting less than full value. The unused or excess capacity goes to waste. In effect, the organization has overpaid for the device unless it uses the device to the point where it is dangerously close to overload—you cannot buy just part of a device.

Moreover, tax benefits that can be realized from the purchase of a device have to be accrued over years of operation, as depreciation of that device/asset. With a paid service (such as cloud), an operational expenditure, the entire payment (perhaps monthly or quarterly) is tax deductible as an expense.

In the cloud, however, the organization is only paying for what it uses (regardless of the number of devices, or fractions of devices, necessary to handle the load) and no more. This is the *metered service* aspect described earlier. As a result, the organization does not overpay for these assets. However, cloud providers do have excess capacity available to be apportioned to cloud customers, so your organization is always in a position to experience increased demand (even dramatic, rapid, and significant demand) and not be overwhelmed (this is the *rapid elasticity* aspect described earlier).

One way an organization can use hosted cloud services is to augment internal, private data center capabilities with managed services during times of increased demand. We refer to this as *cloud bursting*. The organization might have data center assets it owns, but it can't handle the increased demand during times of elevated need (crisis situations, heavy holiday shopping periods, and so on), so it rents the additional capacity as needed from an external cloud provider. See Figure 1.1.

Therefore, with deployment to a cloud environment, the organization realizes cost savings immediately (not paying for unused resources) and avoids a costly risk (the possibility of loss of service due to increased demand).

Cloud Governance

Cloud services can quickly spring up all over an organization as individual business units make adoption decisions without coordinating with the IT department or other business units.

Cloud governance programs try to bring all of an organization's cloud activities under more centralized control. They serve as a screening body helping to ensure that cloud services

used by the organization meet technical, functional, and security requirements. They also provide a centralized point of monitoring for duplicative services, preventing different business units from spending money on similar services when consolidation would reduce both costs and the complexity of the operating environment.

Building a centralized governance program also helps organizations avoid the use of *shadow IT,* where functional units discover and provision cloud services on their own to satisfy unmet technical needs.

FIGURE 1.1 Rapid scalability allows the customer to dictate the volume of resource usage.

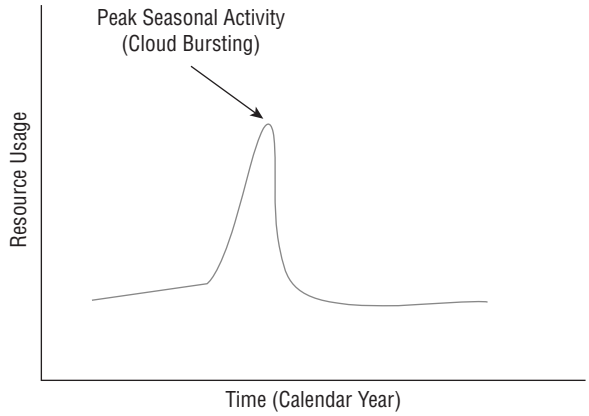

Reduction in Personnel Costs

For most organizations (other than those that deliver IT services), managing data is not a core competency, much less a profitable line of business. Data management is also a specialized skill, and people with IT experience and training are relatively expensive (compared to employees in other departments). The personnel required to fulfill the physical needs of an internal IT environment represent a significant and disproportionally large investment for the organization. In moving to the cloud, the organization can largely divest itself of a large percentage, if not a majority, of these personnel.

Reduction in Operational Costs

Maintaining and administering an internal environment takes a great deal of effort and expense. When an organization moves to the cloud, the cost becomes part of the price of the service, as calculated by the cloud provider. Therefore, costs are lumped in with the flat-rate cost of the contract and will not increase in response to enhanced operations (scheduled updates, emergency response activities, and so on).

Transferring Some Regulatory Costs

Some cloud providers may offer holistic, targeted regulatory compliance packages for their customers. For instance, the cloud provider might have a set of controls that can be applied to a given customer's cloud environment to ensure the mandates of the payment card industry (PCI) are met. Any customer wanting that package can specify so in a service contract instead of trying to delineate individual controls à la carte. In this manner, the cloud customer can decrease some of the effort and expense they might otherwise incur in trying to come up with a control framework for adhering to the relevant regulations.

We will go into more detail about service-level agreements, or service contracts, in later chapters.

While it is possible to transfer some of the responsibilities and costs to service providers or insurance companies, it simply isn't possible to transfer all responsibility to external providers. If your organization collects *personally identifiable information (PII)*, you remain ultimately responsible for any breaches or releases of that data, even if you are using a cloud service and the breach/release results from negligence or attack on the part of the cloud provider. You might be able to transfer some of the financial risk, but you still may be subject to regulatory and/or reputational risk in the wake of a breach.

Reduction in Costs for Data Archival/Backup Services

Offsite backups are standard practice for both long-term data archival and disaster recovery purposes. Having a cloud-based service for this purpose is sensible and cost-efficient even if the organization does not conduct its regular operations in the cloud. However, moving operations into the cloud can create an economy of scale when combined with the archiving/backup usage; this can lead to an overall cost savings for the organization. It can enhance the business continuity and disaster recovery (BC/DR) strategy for the organization as well.

Intended Impact

All of these benefits can be enumerated according to dollar value: each potential cost-saving measure can be quantified. Senior management—with input from subject matter experts—needs to balance the potential financial benefits against the risks of operating in the cloud. It is this cost-benefit calculation, driven by business needs but informed by security concerns, that will allow senior management to decide whether a cloud migration of the organization's operational environment makes sense.

Return on investment (ROI) is a term related to cost-benefit measures. It is used to describe a profitability ratio. It is generally calculated by dividing net profit by net assets.

A great many risks are associated with cloud migration as well. We will be addressing these in detail throughout this book.

Cloud Computing Service Categories

Cloud services are often offered in terms of three general categories, based on what the vendor offers and the customer needs and the responsibilities of each according to the service contract. (ISC)^2 expects you to understand these three models for testing purposes. These categories are *software as a service (SaaS)*, *infrastructure as a service (IaaS)*, and *platform as a service (PaaS)*, as shown in Figure 1.2. In the following sections, we'll review each of them in turn.

FIGURE 1.2 Cloud service categories

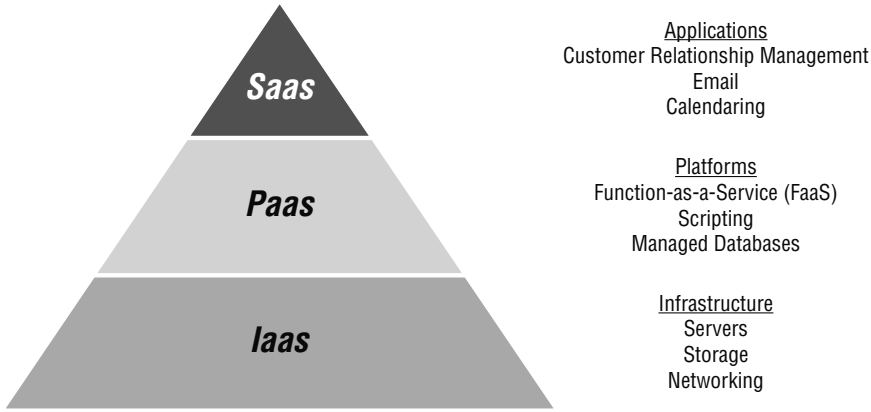

Some vendors and consultants demonstrate a lot of zeal in capitalizing on the popularity of the "cloud" concept and incorporate the word into every term they can think of in order to make their products more appealing. We see a broad proliferation of such labels as networking as a service (NaaS), compliance as a service (CaaS), and data science as a service (DSaaS), but they're mostly just marketing techniques. The only service categories you'll need to know for both the exam and your use as a practitioner are IaaS, PaaS, and SaaS.

Software as a Service

In *software as a service (SaaS)* offerings, the public cloud provider delivers an entire application to its customers. Customers don't need to worry about processing, storage, networking,

or any of the infrastructure details of the cloud service. The vendor writes the application, configures the servers, and basically gets everything running for customers, who then simply use the service. Very often these services are accessed through a standard web browser, so very little, if any, configuration is required on the customer's end.

Common examples of software as a service application capability types include email delivered by Google Apps or Microsoft Office 365 and storage services that facilitate collaboration and synchronization across devices, such as Box and Dropbox. SaaS applications can also be very specialized, such as credit card processing services and travel and expense reporting management.

Infrastructure as a Service

Customers of *infrastructure as a service (IaaS)* vendors purchase basic computing resources from vendors and piece them together to create customized IT solutions. For example, IaaS vendors might provide compute capacity, data storage, and other basic infrastructure building blocks. The four largest vendors in the IaaS space are Amazon Web Services (AWS), Microsoft Azure, Google Compute Engine, and Alibaba.

IaaS includes these common infrastructure capability types:

- Virtualized servers that run on shared hardware
- Block storage that is available as disk volumes
- Object storage that maintains files in buckets
- Networking capacity to connect servers to each other and the internet
- Orchestration capabilities that automate the work of administering cloud infrastructure

IaaS vendors provide on-demand, self-service access to computing resources, allowing customers to request resources when they need them and immediately gain access to them.

Platform as a Service

In the final category of public cloud computing, *platform as a service (PaaS),* vendors provide customers with a platform where they can run their own application code without worrying about server configuration. This is a middle ground between IaaS and SaaS. With PaaS, customers don't need to worry about managing servers but are still able to run their own code.

Function as a service (FaaS) is a common PaaS capability where the customer creates specialized functions that run either on a schedule or in response to events. Examples of FaaS offerings include AWS Lambda, Microsoft Azure Functions, and Google Cloud Functions.

Platform capability types also include cloud-based database engines and services as well as "big data"–style services, such as data warehousing and data mining. The provider offers access to the back-end engine/functionality, while the customer can create/install various apps/APIs to access the back end.

Cloud Deployment Models

Cloud deployment models describe different approaches to the way that organizations might implement cloud services. Essentially, they describe the place where the assets are hosted and who has access to them. The major cloud deployment models are private cloud, public cloud, hybrid cloud, multi-cloud, and community cloud.

Private Cloud

Organizations using the *private cloud* model want to gain the flexibility, scalability, agility, and cost effectiveness of the cloud but don't want to share computing resources with other organizations. In the private cloud approach, the organization builds and runs its own cloud infrastructure or pays another organization to do so on its behalf.

A private cloud is typified by resources dedicated to a single customer; no other customers will share the underlying resources (hardware and perhaps software). Therefore, private clouds are *not* multitenant environments.

Public Cloud

The *public cloud* uses the multitenancy model. In this approach, cloud providers build massive infrastructures in their data centers and then make those resources available to all comers. The same physical hardware may be running workloads for many different customers at the same time.

Hybrid Cloud

Organizations adopting a *hybrid cloud* approach use a combination of public and private cloud computing. In this model, they may use the public cloud for some computing workloads but they also operate their own private cloud for some workloads, often because of data sensitivity concerns.

Multi-Cloud

While many organizations pick a single public cloud provider to serve as their infrastructure partner, some choose to adopt a *multi-cloud* approach that combines resources from two or more public cloud vendors. This approach allows organizations to take advantage of service and price differences, but it comes with the cost of added complexity.

Community Cloud

Community clouds are similar to private clouds in that they are not open to the general public, but they are shared among several or many organizations that are related to each

other in a common community. For example, a group of colleges and universities might get together and create a community cloud that provides shared computing resources for faculty and students at all participating schools.

Gaming communities might be considered community clouds. For instance, the PlayStation network involves many different entities coming together to engage in online gaming: Sony hosts the identity and access management (IAM) tasks for the network, a particular game company might host a set of servers that run information rights management (IRM) functions and processing for a specific game, and individual users conduct some of their own processing and storage locally on their own PlayStations. In this type of community cloud, ownership of the underlying technologies (hardware, software, and so on) is spread throughout the various members of the community.

A community cloud can also be provisioned by a third party on behalf of the various members of the community. For instance, a cloud provider might offer a FedRAMP cloud service, for use only by U.S. federal government customers. Any number of federal agencies might subscribe to this cloud service (say, the Department of Agriculture, Health and Human Services, the Department of the Interior, and so on), and they will all use underlying infrastructure that is dedicated strictly for their use. Any customer that is not a U.S. federal agency will not be allowed to use this service, as nongovernmental entities are not part of this particular community. The cloud provider owns the underlying infrastructure, but it's provisioned and made available solely for the use of the specific community.

Exam Tip

When you take the exam, you should remember that no one cloud deployment model is inherently superior to the others. Organizations may wish to use a public cloud–heavy approach to achieve greater cost savings while others may have regulatory requirements that prohibit the use of shared tenancy computing.

Multitenancy

The public cloud is built upon the operating principle of multitenancy. This simply means that many different customers share use of the same computing resources. The physical servers that support our workloads might be the same as the physical servers supporting your workloads.

In an ideal world, an individual customer should never see the impact of multitenancy. Servers should appear completely independent of each other and enforce the principle of isolation. From a privacy perspective, one customer should never be able to see data belonging to another customer. From a performance perspective, the actions that one customer takes

should never impact the actions of another customer. Preserving isolation is the core crucial security task of a cloud service provider.

Multitenancy allows cloud providers to oversubscribe their resources. Almost all computing workloads vary in their needs over time. One application might have a high CPU utilization for a few hours in the morning, while another uses small peaks throughout the day and others have steady use or different peaks.

Oversubscription means that cloud providers can sell customers a total capacity that exceeds the actual physical capacity of their infrastructure because, in the big picture, customers will never use all of that capacity simultaneously. When we fit those workloads together, their total utilization doesn't ever exceed the total capacity of the environment.

Multitenancy works because of resource pooling. The memory and CPU capacity of the physical environment are shared among many different users and can be reassigned as needed.

Of course, sometimes this concept breaks down. If customers do suddenly have simultaneous demands for resources that exceed the total capacity of the environment, performance degrades. This causes slowdowns and outages. Preventing this situation is one of the key operational tasks of a cloud service provider, and they work hard to manage workload allocation to prevent this from happening.

Cloud Computing Roles and Responsibilities

In the world of cloud computing, people and organizations take on different roles and responsibilities. As a cloud security professional, it's important that you understand these different roles. The two primary roles in the world of cloud computing are those of the cloud service provider and the customer.

The *cloud service provider* is the business that offers cloud computing services for sale to third parties. The cloud service provider is responsible for building and maintaining their service offerings. Cloud service providers may do this by creating their own physical infrastructure, or they might outsource portions of their infrastructure to other cloud service providers. In that case, they are also cloud customers!

Customers are the consumers of cloud computing services. They use cloud services as the infrastructure, platforms, and/or applications that help them run their own businesses.

The relationship between the cloud service provider and the customer varies depending upon the nature, importance, and cost of the service. A customer may never interact with employees at a cloud provider, purchasing services only on a self-service basis, or a cloud provider may have dedicated account representatives who help manage the relationship with different customers.

Cloud service partners play another important role in the cloud ecosystem. These are third-party companies that offer some product or service that interacts with the primary

offerings of a cloud service provider. For example, a cloud service partner might assist a company in implementing a cloud application, or it might offer a security monitoring service that provides operational assistance with using a cloud infrastructure product. Large cloud service providers commonly have a certification program that designates third-party vendors as official partners.

Regulators also play an important role in the cloud ecosystem. Different regulatory agencies may have authority over your business depending upon the locations where your organization does business and the industries in which you operate. Make sure you consult the rules published by different regulators to ensure that your use of cloud computing resources doesn't run afoul of their requirements.

Finally, the last role that we'll discuss is that of the *cloud access security broker (CASB)*. These are cloud service providers who offer a managed identity and access management service to cloud customers that integrates security requirements across cloud services. We'll talk more about CASBs later in this book.

When organizations use public cloud resources, they must understand that security in the public cloud follows a shared responsibility model. Depending upon the nature of the cloud service, the cloud service provider is responsible for some areas of security while the customer is responsible for other areas. For example, if you purchase a cloud storage service, it's your responsibility to know what data you're sending to the service and probably to configure access control policies that say who may access your data. It's the provider's responsibility to encrypt data under their care and correctly implement your access control policies.

Cloud Computing Reference Architecture

The International Organization for Standardization (ISO) publishes a *cloud reference architecture* in its document ISO 17789. This document lays out a common terminology framework that assists cloud service providers, cloud service customers, and cloud service partners in communicating about roles and responsibilities.

Exam Tip

Before we dive in, it's important to note that the cloud reference architecture concepts are a helpful framework but there are no reference architecture police. You should feel free to use whatever terms and framework make sense for your organization. That said, the CCSP exam does explicitly cover the cloud reference architecture, so be sure you're familiar with it.

The reference architecture defines different cloud computing activities that are the responsibility of different organizations in the cloud ecosystem.

Let's begin by talking about the responsibilities of the cloud service customer. The reference architecture says that the following activities are the responsibilities of the customer:

- Use cloud services
- Perform service trials
- Monitor services
- Administer service security
- Provide billing and usage reports
- Handle problem reports
- Administer tenancies
- Perform business administration
- Select and purchase service
- Request audit reports

Cloud service providers have many more responsibilities:

- Prepare systems and provide cloud services
- Monitor and administer services
- Manage assets and inventories
- Provide audit data
- Manage customer relationships and handle customer requests
- Perform peering with other cloud providers
- Ensure compliance
- Provide network connectivity

Finally, cloud service partners have varying activities depending upon the type of partner:

- Design, create, and maintain service components
- Test services
- Perform audits
- Set up legal agreements
- Acquire and assess customers
- Assess the marketplace

In the real world, these activities may shift around depending upon the nature of each organization and the cloud services being provided. However, the reference architecture provides us with a starting point.

Virtualization

The world of enterprise computing has changed dramatically over the years. The advent of virtualization is one of those transformative changes and is the driving force behind cloud computing infrastructure.

It was only a few decades ago that enterprise computing was confined to the world of the data center and its mainframe. Dozens of computing professionals carefully tended to this very valuable resource that served as the organization's electronic nerve center.

Then, in the 1980s and 1990s, the enterprise IT landscape shifted dramatically. We moved away from the world of monolithic mainframes to a new environment of client-server computing. This shift brought tremendous benefits. First, it put computing power right on the desktop, allowing users to perform many actions directly on their machines without requiring mainframe access. Centralized computing improved also, by allowing the use of dedicated servers for specific functions. It became much easier to maintain data centers with discrete servers than to tend to a cranky mainframe.

Over the past decade, we've seen another shift in the computing landscape. The client-server model served us well, but it also resulted in wasted resources. Data center managers realized that most of the time, many of their servers were sitting idle, waiting for a future burst in activity. That's not very efficient. Around that same time *virtualization* technology became available that allows many different virtual servers to make use of the same underlying hardware. This shared hardware platform makes it easy to shift memory, storage, and processing power to wherever it's needed at the time. Virtualization platforms like VMware and Microsoft Hyper-V make this possible.

At a high level, virtualization platforms involve the use of a *host machine* that actually has physical hardware. That hardware then hosts several or many *virtual guest machines* that run operating systems of their own.

Hypervisors

The host machine runs special software known as a *hypervisor* to manage the guest virtual machines (VMs). The hypervisor basically tricks each guest into thinking that it is running on its own hardware when, in reality, it's running on the shared hardware of the host machine. The operating system on each guest machine has no idea that it is virtualized, so software on that guest machine can function in the same way as it would on a physical server.

There are two different types of hypervisors, as shown in Figure 1.3.

In a *Type 1 hypervisor*, also known as a bare metal hypervisor, the hypervisor runs directly on top of the hardware and then hosts guest operating systems on top of that. This is the most common form of virtualization found in data centers.

In a *Type 2 hypervisor*, the physical machine actually runs an operating system of its own and the hypervisor runs as a program on top of that operating system. This type of

virtualization is commonly used on personal computers. Common hypervisors used in this scenario are VirtualBox and Parallels.

FIGURE 1.3 Type 1 and Type 2 hypervisors

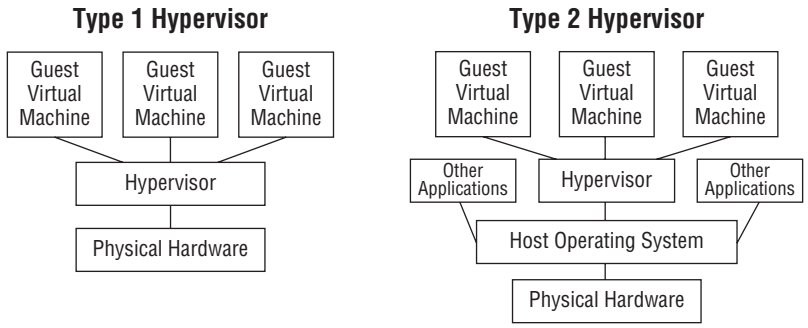

Virtualization Security

From a security perspective, virtualization introduces new concerns around VM *isolation*. In a physical server environment, security teams know that each server runs on its own dedicated processor and memory resources and that if an attacker manages to compromise the machine, they will not have access to the processor and memory used by other systems. In a virtualized environment, this may not be the case if the attacker is able to break out of the virtualized guest operating system. This type of attack is known as a *VM escape* attack.

Virtualization technology is designed to enforce isolation strictly, and the providers of virtualization technology take seriously any vulnerabilities that might allow VM escape. Security professionals working in virtualized environments should pay particular attention to any security updates that affect their virtualization platforms and apply patches promptly.

There's one other security issue associated with virtualization that you should be aware of when preparing for the exam. Virtualization makes it incredibly easy to create new servers in a data center. Administrators can usually create a new server with just a few clicks. While this is a tremendous convenience, it also leads to a situation known as *VM sprawl*, where there are large numbers of unused and abandoned servers on the network. This is not only wasteful, it's also a security risk because those servers are not being properly maintained and may accumulate serious security vulnerabilities over time if they are not properly patched.

One of the major benefits of virtualization is that it allows an organization to adopt an *ephemeral computing* strategy. *Ephemeral* just means temporary, or lasting for a short period of time. Ephemeral computing means that you can create computing resources, such as servers and storage spaces, to solve a particular problem and then get rid of them as soon as you no longer need them. If you tried to do this in the physical world, you'd be installing and removing equipment as needs change. With virtualization, you can create and destroy servers and other computing resources with the click of a button.

Cloud Shared Considerations

As you decide the appropriate ways for your organization to make use of the cloud, you should keep a set of other considerations in mind. Let's take a look at some of the important factors that organizations must consider during cloud deployment efforts.

Security and Privacy Considerations

At a high level, the security and privacy concerns facing cloud computing users are the same as those faced by any cybersecurity professional. We have the three main goals of cybersecurity: confidentiality, integrity, and availability. These three goals, commonly referred to as the CIA triad, are shown in Figure 1.4. We supplement the CIA triad with an additional goal of protecting individual privacy.

FIGURE 1.4 The CIA triad

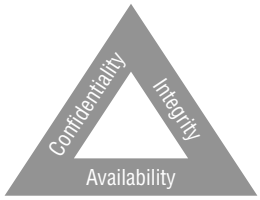

- *Confidentiality* seeks to protect assets (information and systems) from unauthorized access.
- *Integrity* seeks to protect those same assets against unauthorized modification.
- *Availability* seeks to ensure that assets are available for authorized use when needed without disruption.

Privacy adds another dimension to these requirements by ensuring that we respect the rights of confidentiality not only of our own organization but also of the individuals whose personal information we store, process, and transmit. The presence of these goals is nothing new. We worry about the confidentiality, integrity, availability, and privacy of information wherever we use it. The only thing that changes in the world of cloud computing is that we may have more partners to include in our security planning work. This introduces three new concerns: governance, auditability, and regulatory oversight.

Cloud computing *governance* efforts help an organization work through existing and planned cloud relationships to ensure that they comply with security, legal, business, and other constraints. Most organizations set up a cloud governance structure designed to vet potential vendors, manage relationships, and oversee cloud operations. These governance structures are crucial to organizing cloud efforts and ensuring effective oversight.

Auditability is an important component of governance. Cloud computing contracts should specify that the customer has the right to audit cloud providers, either directly or through a third party. These audits may take place on a scheduled or unplanned basis, allowing the customer to gain assurance that the cloud vendor is meeting its security obligations. The audits may also include operational and financial considerations.

Finally, *regulatory oversight* exists in the cloud world, just as it does in the realm of on-premises computing. Organizations subject to HIPAA, FERPA, PCI DSS, or other cybersecurity regulations must ensure that cloud providers support their ability to remain compliant with those obligations. Some regulatory schemes include specific provisions about how organizations ensure that third-party providers remain compliant, such as using only certified providers or requiring written agreements with providers that the provider's handling of data will be consistent with regulations.

Operational Considerations

Just as the cloud raises security considerations that organizations must take into account through their governance structures, it also raises operational concerns. As with the security concerns we discussed, the operational considerations for cloud computing are quite similar to those that we encounter during on-premises operations. Let's look at a few of these considerations.

Availability and Performance

First, we mentioned availability as a security consideration, but it is also an operational consideration. One of the core measures of cloud performance is that service's availability. What percentage of the time is the service up and running and meeting customer needs?

We can increase availability by increasing our *resiliency*. Resiliency is the ability of the cloud infrastructure to withstand disruptive events. For example, we can use redundant servers to protect against the failure of a single server and we can use multiple cloud data centers to protect against the failure of an entire data center.

Performance is a closely related concept. How well can the cloud service stand up to the demands that we place on it? If we encounter an extremely busy period, will the service continue to respond at an appropriate rate?

All three of these considerations: availability, resiliency, and performance, are crucial issues to cloud operations. Customers should negotiate specific service levels with vendors during the contracting process and then document those service levels in written agreements called *service-level agreements (SLAs)*. SLAs specify the requirements that the vendor agrees to meet and commonly include financial penalties if the vendor fails to live up to operational obligations.

Maintenance and Version Control

IT teams around the world also know the importance of scheduled maintenance and *version control*. Managing change is a difficult issue in enterprise IT, and those concerns don't go away in the cloud. In fact, they become more complex because IT teams must not only

coordinate their own maintenance schedules but also consider the maintenance schedules of cloud providers. Does the provider have scheduled maintenance periods? If so, IT teams must consider how those periods will impact business operations.

Version control allows organizations to manage the development of software by tracking different versions being worked on by different developers. It is an essential component of any software development program. Version control may also be used to track the configuration of systems and applications.

Outsourcing Issues

Moving to the cloud also introduces some cloud-specific operational considerations that come as a result of outsourcing parts of our IT operations. Let's talk about three of these concerns.

First, organizations moving to the cloud or between cloud vendors should consider the importance of *reversibility*. What if something goes wrong operationally, technically, or financially? How difficult would it be to restore the original operations and reverse the move? Organizations should make rollback plans part of every transition plan.

Similarly, organizations should strive to avoid *vendor lock-in* whenever possible. Portability is a design principle that says workloads should be designed so that they don't leverage vendor-specific features and may be more easily shifted between cloud providers. This isn't always possible, but it is a good design practice.

Each vendor relationship should also provide the ability to export data when required. The vendor relationship will eventually come to an end and the organization will need the ability to retrieve any business data stored with the cloud provider to support a transition to another provider or data archive.

Finally, organizations should consider the *interoperability* of cloud providers whenever adopting a new cloud solution. This is especially important for SaaS and PaaS products. IT teams are called upon to integrate solutions regularly, and the ability of a vendor to support those integrations is crucial. Imagine the impact if your expense reporting system couldn't interoperate with your financial system, or if your storage provider didn't interoperate with your web content management solution. Interoperability is crucial.

Emerging Technologies

As you prepare for the CCSP exam, you need to be familiar with a set of emerging technologies that are especially significant in the world of cloud computing.

Machine Learning and Artificial Intelligence

Machine learning is a technical discipline designed to apply the principles of data science and statistics to uncover knowledge hidden in the data that we accumulate every day. Machine learning techniques analyze data to uncover trends, categorize records, and help us run our businesses more efficiently.

Machine learning is a subset of a broader field called *artificial intelligence (AI)*. AI is a collection of techniques, including machine learning, that are designed to mimic human thought processes in computers, at least to some extent.

As we conduct machine learning, we have a few possible goals:

- *Descriptive analytics* simply seeks to describe our data. For example, if we perform descriptive analytics on our customer records, we might ask questions like, what proportion of our customers are female? And how many of them are repeat customers?
- *Predictive analytics* seek to use our existing data to predict future events. For example, if we have a dataset on how our customers respond to direct mail, we might use that dataset to build a model that predicts how individual customers will respond to a specific future mailing. That might help us tweak that mailing to improve the response rate by changing the day we send it, altering the content of the message, or even making seemingly minor changes like altering the font size or paper color.
- *Prescriptive analytics* seek to optimize our behavior by simulating many scenarios. For example, if we want to determine the best way to allocate our marketing dollars, we might run different simulations of consumer response and then use algorithms to prescribe our behavior in that context. Similarly, we might use prescriptive analytics to optimize the performance of an automated manufacturing process.

Cloud computing has revolutionized the world of machine learning. Many of the applications where we apply artificial intelligence techniques today simply wouldn't have been possible with the scalable, on-demand computing offered by the cloud. Cloud providers now offer very specialized services designed to help organizations design, build, and implement machine learning models.

Blockchain

The second emerging technology that you'll need to understand for the CCSP exam is *blockchain* technology. The blockchain is, in its simplest description, a distributed immutable ledger. This means that it can store records in a way that distributes those records among many different systems located around the world and do so in manner that prevents anyone from tampering with the records. The blockchain creates a data store that nobody can tamper with or destroy.

The first major application of the blockchain is *cryptocurrency*. The blockchain was originally invented as a foundational technology for Bitcoin, allowing the tracking of Bitcoin transactions without the use of a centralized authority. In this manner, blockchain allows the existence of a currency that has no central regulator. Authority for Bitcoin transactions is distributed among all participants in the Bitcoin blockchain. While cryptocurrency is the blockchain application that has received the most attention, there are many other uses for a distributed immutable ledger, so much so that new applications of blockchain technology seem to be appearing every day, as in the following examples:

- Property ownership records could benefit tremendously from a blockchain application. This approach would place those records in a transparent, public repository that is protected against intentional or accidental damage.

- Blockchain might also be used to track supply chains, providing consumers with confidence that their produce came from reputable sources and allowing regulators to easily track down the origin of recalled produce.
- Blockchain applications can track vital records, such as passports, birth certificates, and death certificates. The possibilities are endless.

Cloud computing enables blockchain by providing computing resources that are scalable, economically efficient, and globally distributed.

Internet of Things

The *Internet of Things (IoT)* is the third emerging technology covered on the CCSP exam. *IoT* is a term used to describe connecting nontraditional devices to the internet for data collection, analysis, and control. We see IoT applications arising in the home and workplace.

On the home front, it's hard to walk around your house or the local consumer electronics store without seeing a huge number of devices that are now called "smart this" or "smart that." We can now have in our homes a smart television, a smart garage door, and even a smart sprinkler system. All that *smart* means is that the devices are computer controlled and network connected.

IoT technology began by taking some of the more common computing devices in our homes, such as game consoles and printers, and making them first network connected and then wireless. Manufacturers quickly realized that we wanted connectivity to enable multiplayer games and printing without cables and then brought this technology into the home. From there, the possibilities were endless and wireless smart devices spread throughout the home, and even into the garage with the advent of smart cars that expose in-vehicle computing systems to the drivers and passengers.

All of these devices come with security challenges as well. First, it is often difficult for the consumer to update the software on these devices. While the device may run slimmed-down versions of traditional operating systems, they don't always have displays or keyboards, so we don't see that our so-called "smart" device is actually running an outdated copy of Windows 95!

Second, these devices connect to the same wireless networks that we use for personal productivity. If a smart device is compromised, it can be a gateway to the rest of our network.

Finally, smart devices often connect back to cloud services for command and control, creating a potential pathway onto our network for external attackers that bypasses the firewall.

Containers

Containers are the next evolution of virtualization. They're a lightweight way to package up an entire application and make it portable so that it can easily move between hardware platforms.

In traditional virtualization, we have hardware that supports a hypervisor and then that hypervisor supports guest virtual machines. Each of those guest machines runs its own

operating system and applications, allowing the applications to function somewhat independently of the hardware. You can move a virtual machine from hardware to hardware, as long as the machines are running the same hypervisor. One of the downsides to traditional virtualization is that virtual machines are somewhat heavy. Each one has to have its own operating system and components. If you're running 10 different Windows virtual servers on a hypervisor, you have the overhead of running 10 different copies of Windows at the same time.

Containerization seeks to reduce this burden by building more lightweight packages. Containers package up application code in a standardized format so that it can be easily shifted between systems. Instead of running a hypervisor, systems supporting containers run a containerization platform. This platform provides a standard interface to the operating system that allows containers to function regardless of the operating system and hardware. The major benefit of containers over virtual machines is that they don't have their own operating systems kernel. The containerization platform allows them to use the host's operating system kernel.

From a security perspective, containers share many of the same considerations as virtualized systems. The containerization platform must strictly enforce isolation to ensure that containers cannot access the data or resources allocated to other containers. As long as this isolation remains intact, containers are a highly secure option for lightweight virtualized computing.

Quantum Computing

Quantum computing is an area of advanced theoretical research in computer science and physics. The theory is that we can use principles of quantum mechanics to replace the binary 1 and 0 bits of digital computing with multidimensional quantum bits known as qubits.

Exam Tip

Quantum computing is one of the emerging technologies that you need to be familiar with for the CCSP exam, but honestly, there's not much that you need to know. That's because, as of today, quantum computers are confined to theoretical research. Nobody has yet developed a practical implementation of a useful quantum computer.

If practical quantum computers do come on the scene, they have the potential to revolutionize the world of computer science by providing the technological foundation for the most powerful computers ever developed. Those computers would quickly upend many of the principles of modern cybersecurity. For example, it is possible that a quantum computer could render all modern cryptography completely ineffective and require the redesign of new, stronger quantum cryptography algorithms. But that's all theory for now. Unless you're a

research physicist, there won't be much impact of quantum computing on your world for the foreseeable future.

Edge and Fog Computing

The emergence of the Internet of Things is also dramatically changing the way that we provision and use computing. We see the most dramatic examples of the Internet of Things in our everyday lives, from connected and semiautonomous vehicles to smart home devices that improve the way we live and travel.

However, many of the applications of the Internet of Things occur out of sight. Industrial applications of IoT are transforming manufacturing. We're seeing the rise of microsatellites that bring scientific instruments and other devices into Earth orbit and beyond. Even agriculture is changing dramatically with the sensors, information, and analytics that IoT makes possible.

The cloud computing model doesn't always fit these applications. When your sensors are far away from cloud data centers and either not connected to a network or connected by very low bandwidth connections, the model starts to break down. It simply doesn't make sense to transfer all of the data back to the cloud and have it processed there.

Edge computing is an approach that brings many of the advances of the cloud to the edge of our networks. It involves placing processing power directly on remote sensors and allowing them to perform the heavy lifting required to process data before transmitting a small subset of that data back to the cloud.

Fog computing is a related concept that involves placing gateway devices out in the field to collect information from sensors and perform that correlation centrally, but still at the remote location, before returning data to the cloud. Together, edge and fog computing promise to increase our ability to connect IoT devices and the cloud.

Confidential Computing

Confidential computing is a new and emerging focus for organizations operating in extremely secure environments, such as the military and defense sector. It extends security throughout the entire computing process.

Let's take a look at what this means. In a normal client-server computing model, we have a client that might want to request that a server take some action on the client's behalf. That server might have to access some data that is maintained in a separate storage environment and then perform some processing on that data.

In a traditional computing model, we know that we need to add some security to this process. We apply encryption to the data that is kept on the storage devices so that unauthorized individuals can't access it. That's called protecting data at rest. We also add encryption to network communications to prevent eavesdropping attacks between the client and the server and between the server and storage. That's called protecting data in motion.

However, we generally don't worry about the data that's being actively processed by the server. We don't apply encryption to data in memory, which is known as data in use.

Confidential computing adds protection for code and data in memory. It does this by offering *trusted execution environments (TEEs)*. These trusted environments guarantee that no outside process can view or alter the data being handled within the environment. That provides an added assurance that data is safe through all stages of the computing lifecycle.

DevOps and DevSecOps

IT organizations around the world are quickly embracing the *DevOps* philosophy to improve the interactions between software developers and technology operations teams. The DevOps movement seeks to combine two worlds that have often found themselves in conflict in the past.

Software developers are charged with creating code: building applications and integrations that meet the needs of customers and the business. They are motivated to rapidly release code and meet those demands.

IT operations staff are charged with maintaining the infrastructure and keeping the enterprise stable. They are often wary of change because change brings the possibility of instability. This makes them nervous when developers seek to rapidly deploy new code.

The DevOps movement seeks to bring these two disciplines together in a partnership. DevOps seeks to build collaborative relationships between developers and operators with open communication. The DevOps movement embraces automation as an enabler of both development and operations. DevOps practitioners seek to create the environment where developers can rapidly release new code, while operations staff can provide a stable environment.

The DevOps philosophy is often tightly linked to the *agile software development methodology*. While they are two different concepts, they are deeply related to each other. Developers following these strategies seek to implement a continuous integration software development approach where they can quickly release software updates, creating multiple software releases each day, sometimes even releasing hundreds of updates in a single day.

Cloud computing is one of the enabling technologies for DevOps environments. Specifically, DevOps shops embrace a concept known as *infrastructure as code (IaC)*. In this approach, operations teams no longer manually configure servers and other infrastructure components by logging in and modifying their configurations directly. Instead, they write scripts that specify how to start with a baseline configuration image and then customize it to meet the specific requirements of the situation. For example, an organization might have a standard baseline for a Linux system. When someone needs a new server, they write a script that starts a server using the baseline configuration and then automatically configures it to meet the specific functional needs.

Infrastructure as code separates server configuration from specific physical or virtual servers. This has some clear advantages for the organization.

- **IaC enables scalability**. If the organization needs more servers, the code can create as many as necessary extremely rapidly.
- **IaC reduces user error through the use of immutable servers**. This means that engineers don't ever log into or modify servers directly. If they need to make a change, they modify the code and create new servers.
- **IaC makes testing easy**. Developers can write code for new servers and spin up a fully functional test environment without affecting production. Once they've verified that the new code works properly, they can move it to production and destroy the old servers.

The DevOps approach to IT provides many different benefits to the organization. Security teams can also benefit from this approach by using security automation techniques. There's no reason that cybersecurity teams can't embrace the DevOps philosophy and build security infrastructure and analysis tools using an infrastructure as code approach.

When DevOps is used in a cybersecurity program, it is often referred to as *DevSecOps* and introduces a "security as code" approach to cybersecurity. As organizations move to DevOps strategies, cybersecurity teams will need to evolve their practices to provide value in this new operating environment.

Summary

In this chapter, we have examined business requirements, cloud definitions, cloud computing roles and responsibilities, and foundational concepts of cloud computing. This chapter has provided an introductory foundation for these topics. We will explore each of them in more detail as we move ahead.

Exam Essentials

Explain the different roles in cloud computing. Cloud service providers offer cloud computing services for sale. Cloud service customers purchase these services and use them to meet their business objectives. Cloud service partners assist cloud service customers in implementing the services they purchase from providers. Cloud access service brokers offer an intermediary layer of security for cloud users. Regulators define requirements for operating in the cloud with sensitive data.

Identify the key characteristics of cloud computing. The key characteristics of cloud computing are on-demand self-service, broad network access, multitenancy, rapid elasticity and scalability, resource pooling, and measured service.

Explain the three cloud service categories. The three cloud service categories are software as a service (SaaS), infrastructure as a service (IaaS), and platform as a service (PaaS). In a SaaS offering, the provider runs a complete application in the cloud that is sold to customers.

In IaaS offerings, the provider sells technology building blocks to customers, who assemble their own solutions. In PaaS offerings, the provider sells an environment where customers may run their own code.

Describe the five cloud deployment models. Public cloud deployments use multitenancy to provide services to many customers on shared hardware. Private cloud environments use hardware that is dedicated to a single customer. Hybrid cloud environments make use of both public and private cloud services. Community clouds are dedicated to a group of customers with a shared characteristic. Some organizations choose to combine cloud services from several providers in a multi-cloud deployment.

Identify important related technologies. Cloud computing benefits from and serves several related technologies. These include data science, machine learning, artificial intelligence, blockchain, the Internet of Things, containers, quantum computing, edge computing, fog computing, confidential computing, and DevSecOps.

Explain the shared considerations in the cloud. As organizations decide whether to use cloud services, they must analyze several important considerations. These include interoperability, portability, reversibility, availability, security, privacy, resiliency, performance, governance, maintenance and versioning, service levels and service-level agreements, auditability, regulatory concerns, and the impact of outsourcing.

Review Questions

You can find the answers to the review questions in Appendix A.

1. Which of the following is *not* a common cloud service model?
   - **A.** Software as a service (SaaS)
   - **B.** Programming as a service (PaaS)
   - **C.** Infrastructure as a service (IaaS)
   - **D.** Platform as a service (PaaS)
2. Which one of the following emerging technologies, if fully implemented, would jeopardize the security of current encryption technology?
   - **A.** Quantum computing
   - **B.** Blockchain
   - **C.** Internet of Things
   - **D.** Confidential computing
3. Cloud vendors are held to contractual obligations with specified metrics by _____.
   - **A.** Service-level agreements (SLAs)
   - **B.** Regulations
   - **C.** Law
   - **D.** Discipline
4. _____ drive security decisions.
   - **A.** Customer service responses
   - **B.** Surveys
   - **C.** Business requirements
   - **D.** Public opinion
5. If a cloud customer cannot get access to the cloud provider, this affects what portion of the CIA triad?
   - **A.** Integrity
   - **B.** Authentication
   - **C.** Confidentiality
   - **D.** Availability

6. You recently worked with a third-party vendor to help you implement a SaaS offering provided by a different company. Which one of the following cloud service roles is not represented here?
   - **A.** Regulator
   - **B.** Customer
   - **C.** Provider
   - **D.** Partner

7. Which of the following hypervisor types is most likely to be seen in a cloud provider's data center?
   - **A.** Type 1
   - **B.** Type 2
   - **C.** Type 3
   - **D.** Type 4

8. All of these are reasons an organization may want to consider cloud migration except _____.
   - **A.** Reduced personnel costs
   - **B.** Elimination of risks
   - **C.** Reduced operational expenses
   - **D.** Increased efficiency

9. The generally accepted definition of cloud computing includes all of the following characteristics except _____.
   - **A.** On-demand self-service
   - **B.** Negating the need for backups
   - **C.** Resource pooling
   - **D.** Measured or metered service

10. You are working on a governance project designed to make sure the different cloud services used in your organization work well together. What goal are you attempting to achieve?
    - **A.** Performance
    - **B.** Resiliency
    - **C.** Reversibility
    - **D.** Interoperability

11. The risk that a customer might not be able to switch cloud providers at a later date is known as _____.
    - **A.** Vendor closure
    - **B.** Vendor lock-out
    - **C.** Vendor lock-in
    - **D.** Vendor synchronization

12. All of these are characteristics of cloud computing except _____.

A. Broad network access

B. Diminished elasticity

C. Rapid scaling

D. On-demand self-service

13. When a cloud customer uploads personally identifiable information (PII) to a cloud provider, who is ultimately responsible for the security of that PII?

A. Cloud provider

B. Regulators

C. Cloud customer

D. The individuals who are the subjects of the PII

14. We use which of the following to determine the critical paths, processes, and assets of an organization?

A. Business requirements

B. Business impact analysis (BIA)

C. Risk Management Framework (RMF)

D. Confidentiality, integrity, availability (CIA) triad

15. If an organization owns all of the hardware and infrastructure of a cloud data center that is used only by members of that organization, which cloud deployment model would this be?

A. Private

B. Public

C. Hybrid

D. Motive

16. The cloud deployment model that features ownership by a cloud provider, with services offered to anyone who wants to subscribe, is known as _____.

A. Private

B. Public

C. Hybrid

D. Latent

17. The cloud deployment model that features joint ownership of assets among an affinity group is known as _____.

A. Private

B. Public

C. Hybrid

D. Community

18. You are concerned that an attacker might be able to use a guest virtual machine to gain access to the underlying hypervisor. What term describes this threat?

A. VM escape

B. SQL injection

C. Man-in-the-middle

D. VM sprawl

19. You are considering purchasing an e-commerce system where the cloud provider runs a hosted application on their own servers. What cloud service category is the provider offering?

A. IaaS

B. PaaS

C. SaaS

D. FaaS

20. If a cloud customer wants to build their own computing environment using storage, networking, and compute resources offered by a cloud provider, which cloud service category would probably be best?

A. IaaS

B. PaaS

C. SaaS

D. FaaS