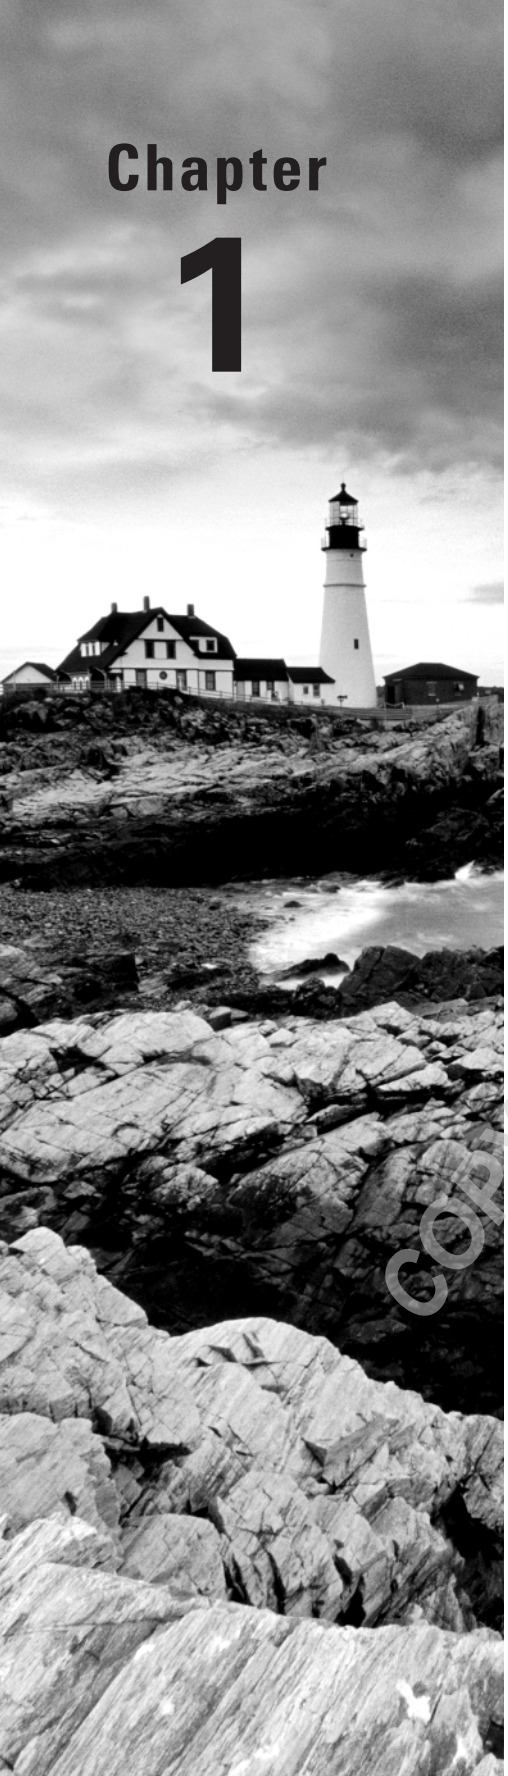

Chapter 1

Architectural Concepts

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

✓ **Domain 1: Architectural Concepts and Design Requirements**

- A. Understand Cloud Computing Concepts
  - A.1 Cloud Computing Definitions
  - A.2 Cloud Computing Roles
  - A.3 Key Cloud Computing Characteristics
  - A.4 Building Block Technologies
- B. Describe Cloud Reference Architecture
  - B.1 Cloud Computing Activities
  - B.2 Cloud Service Capabilities
  - B.3 Cloud Service Categories
  - B.4 Cloud Deployment Models
  - B.5 Cloud Cross-Cutting Aspects
- D. Understanding Design Principles of Secure Cloud Computing
  - D.3 Cost/Benefit Analysis

✓ **Domain 3: Cloud Platform and Infrastructure Security**

- D. Plan Disaster Recovery and Business Continuity Management
  - D.1 Understanding the Cloud Environment
  - D.2 Understanding the Business Requirements

✓ **Domain 6: Legal and Compliance**

- B. Understand Privacy Issues, Including Jurisdictional Variation
  - B.3 Difference Among Confidentiality, Integrity, Availability, and Privacy

This chapter is the foundation for all the other chapters in this study guide. You may find it useful to review this material before reading other chapters.

The CCSP is not a certification of basic computer skills or training; it is a professional certification for practitioners with some background in the field. (ISC)^2 expects that those who want to earn this particular certification already have experience in the industry, have been employed in an InfoSec position in some professional capacity, and have a thorough understanding of many basic areas related to computers, security, business, risk, and networking. Many people taking the test already have other certifications that validate their knowledge and experience such as the CISSP. Therefore, this book will not contain many of the basics that, while testable, you are already expected to know. If you aren't coming from a CISSP background, it would be good to supplement your knowledge with CISSP-focused materials as well.

However, the CCSP Common Body of Knowledge (CBK) contains terminology and concepts that may be expressed in specific ways to include perspectives and usages that may be unique to the CCSP and different from what you are used to dealing with in your normal operations. This chapter is therefore intended as a guide, laying down the foundation for understanding the rest of the material and the CBK as a whole.

Cloud computing has come to mean many things, but the following characteristics have become part of the generally accepted definition:

- Broad network access
- On-demand services
- Resource pooling
- 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."

You can expect to see mention of each of these throughout this book, the CBK, and the exam.

Broad network access means that there should never be network bandwidth bottlenecks. This is generally accomplished with the use of such technologies as advanced routing techniques, load balancers, multisite hosting, and other technologies.

On-demand services refer 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. The services happen in real time.

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).

Finally, measured or metered service simply means that the customer is charged for only what they use and nothing more. This is much like how a water or power company might charge you each month for the services used.

The ISO/IEC standard that provides an overview and vocabulary for cloud computing (ISO/IEC 17788, `www.iso.org/iso/catalogue_detail?csnumber=60544`) includes these traits, and it also adds the characteristic of multitenancy. While it is true that multitenancy is quite often an aspect of most cloud service offerings, it is not exactly a defining element of the field. There are cloud services that do not include multitenancy, as customers can purchase, rent/lease, and stand-alone resources.

Rest assured—we will be going into more detail regarding all of these concepts in the chapters to come.

Real World Scenario

Online Shopping

Think of retail demand during the pre-holiday crush 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 this 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.

It is a great business model, which is why some people say that cloud computing is not a technology but rather a business enabler.

Business Requirements

The IT department is not a profit center; it provides a support function. This is even truer for the security department. Security activities actually 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/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. For example:

- 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 and why and how and where to better understand the nature and intensity of 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.

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.

Many organizations are currently considering moving their network operations to a cloud-based motif. This is not a decision made lightly, and the business requirements must be supported by this transition. As described in the previous paragraphs, there are also different service and delivery models of cloud computing, and an organization must decide which one will optimize success.

Existing State

In this initial effort, a true evaluation and understanding of the business processes, assets, and requirements is 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 the migration.

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.

Possible methods for gathering business requirements include

- Interviewing functional managers
- Interviewing users
- Interviewing senior management
- Surveying customers
- Collecting network traffic
- Inventorying assets
- Collecting financial records
- Collecting insurance records
- Marketing data collection
- 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 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 patents, trademarks, copyrights, and business methodologies.

Quantifying Benefits and Opportunity Cost

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.

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

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 no value. The unused or excess capacity goes to waste. In effect, the organization has overpaid for the device unless the organization uses the device to the point where it is dangerously close to overload—you just cannot buy part of a device.

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* model 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.

One way an organization can use hosted cloud services is to augment internal, private datacenter capabilities with managed services during times of increased demand. We refer to this as "cloud bursting." The organization might have datacenter 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.

FIGURE 1.1 On-demand scalability allows the customer to dictate the volume of resource usage

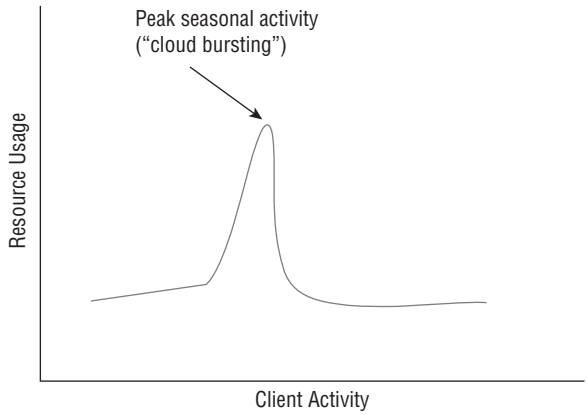

Therefore, with deployment to a cloud environment, the organization realizes cost savings immediately (not paying for unused resources) and avoids a costly risk (the possibility of loss of service due to increased demand).

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 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 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 a 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.

It is, however, crucial to note here (and we'll repeat it throughout the book) that under current laws, no cloud customer can transfer risk or liability associated with the inadvertent or malicious disclosure of personally identifiable information (PII). This is very, very important to understand: your organization, if it holds PII of any kind, is 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. Legally and financially, in the eyes of the court, your organization is always responsible for any unplanned release of PII.

PII is a major component of regulatory compliance, whether the regulation comes in the form of statutes or contractual obligation. Protection of PII will be a large part of our security concern in the cloud.

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. As we'll discuss later in the book, this can enhance the business continuity/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.

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

Cloud Evolution, Vernacular, and Definitions

The arrival of the cloud and its related technology has provided a lot of advantages. To incorporate the cloud and these advantages, it is necessary to understand new terminology and how it relates to the terminology of traditional models. This new technology and its terminology are an integral part of understanding cloud computing service models and cloud computing deployment models.

New Technology, New Options

Fifteen, or even ten years ago, suggesting that organizations hand off their data and operations to a third party that is geographically distant and run by people that most managers in the organization will never meet would have seemed absurd, especially from a security perspective. The risk would have been seen as insurmountable, and ceding that level of control to an outside vendor would have been daunting. Today, a combination of technological capabilities and contractual trust make cloud computing not only appealing but almost a foregone conclusion, in terms of financial viability.

There are specific characteristics that are emblematic of cloud computing. We're going to define them here and offer examples of how each might be demonstrated.

- **Elasticity:** This is the flexibility of allocating resources as needed for immediate usage, instead of purchasing resources according to other variables. For instance, a traditional organization might purchase one desktop for every employee. In that model, the organization would be paying for the entire capacity of the desktop computer—its processing power, its storage capacity, etc.—even though individual users would probably not be using the full capacity of each device at all times.

  In the cloud environment, the organization is paying not for a device, but for the use of a service, when it is being used. The ability of the cloud vendor to offer this type of service (while remaining profitable) is based on the elasticity and the flexibility offered by recent enhancements in technology, including virtualization (we will discuss virtualization further in upcoming chapters). With virtualization, the cloud provider can allocate partial usage of each resource to every user and customer, when those users and customers require it, and nothing more, thereby avoiding wasted, underutilized resources and excess, nonproductive costs.

  In a virtualized environment, users can also access their data from almost any device or platform, and almost any location. This allows portability, availability, and accessibility that exceed previous enterprise environments.
- **Simplicity:** Usage and administration of cloud services ought to be transparent to cloud customers and users; from their perspective, a digital data service is paid for and can be used, with very little additional input other than what is necessary to perform their duties. Proper cloud implementations should not require constant or even frequent interaction between the cloud provider and cloud customer.
- **Scalability:** The organization's computing needs won't remain static: there will be new (and hopefully more) users, customers, and data as the organization continually matures. A cloud service can easily meet those needs, either temporarily or long-term, in a much more cost-efficient manner than a traditional environment, because new computing resources can be assigned and allocated without any significant additional capital investment on the part of the cloud provider, and at an incremental cost to the cloud customer.

The Difference between a Cloud Customer and a Cloud User

A *cloud customer* is anyone who is purchasing a cloud service (which could be an individual or a company), whereas a *cloud user* is just someone using cloud services. It could be an employee of a company who is a cloud customer or just a private individual.

For instance, Company A purchases SaaS services from Cloud Provider X. Company A is a cloud customer. All employees of Company A are cloud users, because they're using the cloud services their employer, a cloud customer, has purchased for their usage.

Not all cloud users are staff of cloud customers, though. Many cloud users are simply individuals who are using publicly available cloud services for their personal purposes, such as a person who has a Gmail account or someone who syncs their smartphone to a free online backup service.

Cloud Computing Service Models

Cloud services are usually offered in terms of three general models, based on what the vendor offers and the customer needs, and the responsibilities of each according to the service contract. These models are Infrastructure as a Service (IaaS), Platform as a Service (PaaS), and Software as a Service (SaaS), as shown in Figure 1.2. In this section, we'll review each of them in turn.

FIGURE 1.2 Cloud Service Models

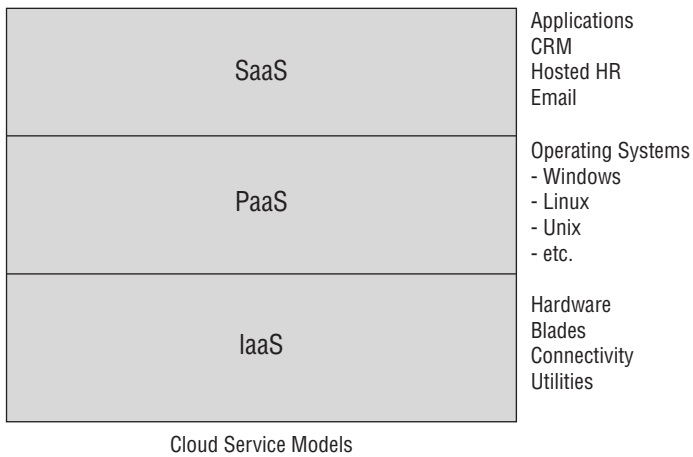

Some vendors and consultants demonstrate a lot of zeal in capitalizing on the popularity of the "cloud" concept and incorporate the word into every term they can think of in order to make their products more appealing. We see a broad proliferation of such labels as Networking as a Service (NaaS), Compliance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

Infrastructure as a Service (IaaS)

The most basic of cloud service offerings, IaaS allows the customer to install all software, including operating systems (OSs) on hardware housed and connected by the cloud vendor.

In this model, the cloud provider has a datacenter with racks and machines and cables and utilities, and administers all these things. However, all logical resources such as software are the responsibility of the customer.

In traditional terms, we might think of this as what used to be considered a "warm site" for BC/DR purposes: the physical space exists, the connectivity exists, and it is available for the customer organization to fill with any type of baseline configuration and populate with any data the customer requires.

IaaS might be optimum for organizations that want enhanced control over the security of their data, or are looking to the cloud for a limited purpose, such as BC/DR or archiving.

Some examples of IaaS would include datacenters that offer Infrastructure as a Service, allowing clients to load whatever operating system and applications they choose. The cloud provider simply supplies the compute, storage, and networking functions.

Platform as a Service (PaaS)

PaaS contains everything included in IaaS, with the addition of OSs. The cloud vendor usually offers a selection of OSs, so that the customer can use any or all of the available choices. The vendor will be responsible for patching, administering, and updating the OS as necessary, and the customer can install any software they deem useful.

This model is especially useful for software development operations (DevOps), as the customer can test their software in an isolated environment without risk of damaging production capabilities, and determine the viability of the software across a range of OS platforms.

Some examples of PaaS include hosting providers that offer not only infrastructure but systems already loaded with a hardened operating system such as Windows Server or Linux.

Software as a Service (SaaS)

SaaS includes everything listed in the previous two models, with the addition of software programs. The cloud vendor becomes responsible for administering, patching, and updating this software as well. The cloud customer is basically only involved in uploading and processing data on a full production environment hosted by the provider.

There are many examples of SaaS configurations, ranging across a spectrum of functionality. Google Docs, Microsoft's Office 365, and QuickBooks Online are all examples of SaaS products.

Some examples of SaaS would include things like customer relationship manager (CRM) software or accounting software hosted in the cloud. The provider takes care of all the infrastructure, compute, and storage needs as well as providing the underlying operating systems and the application itself. All of this is completely transparent to the end user who only sees the application they have purchased.

Cloud Deployment Models

In addition to viewing cloud offerings in terms of what levels of service are involved, another perspective has to do with ownership. You'll be expected to know the facets of both sets of models.

Public

The public cloud is what we typically think of when discussing cloud providers. The resources (hardware, software, facilities, and staff) are owned and operated by a vendor and sold, leased, or rented to anyone (offered to the public—hence the name).

Examples of public cloud vendors include Rackspace, Microsoft's Azure, and Amazon Web Services (AWS).

Private

Private clouds are owned and operated by independent organizations, for the exclusive use of their customers and users. A private cloud can be thought of as the traditional legacy IT environment, with connections made via the web and with remote access capabilities. If your organization hosts a web server and allows access via remote services, that can be considered a private cloud instance.

Examples of private clouds include such things as what used to be called intranets. These often host shared internal applications, storage, and compute resources. One example is an internally hosted SharePoint site.

The terms "public" and "private" can be confusing, because we might think of them in the context of who is offering them instead of who is using them. Remember: A public cloud is owned by a specific company and is offered to anyone who contracts its provided services, whereas a private cloud is owned by a specific organization but is only available to users authorized by that organization.

Community

A community cloud features infrastructure and processing owned and operated by an affinity group; disparate pieces might be owned or controlled by individuals or distinct organizations, but they come together in some fashion to perform joint tasks and functions.

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 digital rights management (DRM) functions and processing for a specific game, and individual users conduct some of their own processing and storage locally on their own PlayStations.

Hybrid

A hybrid cloud, of course, contains elements of the other models. For instance, an organization might want to retain some private cloud resources (say, their legacy production environment, which is accessed remotely by their users), but also lease some public cloud space as well (maybe a PaaS function for DevOps testing, away from the production environment so that there is much less risk of crashing systems in operation).

An example of a hybrid cloud environment might include a hosted internal cloud such as a SharePoint site with a portion carved out for external partners who need to access a shared service. To them it would appear as an external cloud; therefore, it would be operating as a hybrid.

Cloud Computing Roles and Responsibilities

Various entities are involved in cloud service arrangements:

Cloud Service Provider (CSP) The vendor offering cloud services. The CSP will own the datacenter, employ the staff, own and manage the resources (hardware and software), monitor service provision and security, and provide administrative assistance for the customer and the customer's data and processing needs. Examples include Amazon Web Services, Rackspace, and Microsoft's Azure.

Cloud Customer The organization purchasing, leasing, or renting cloud services.

Cloud Access Security Broker (CASB) A third-party entity offering independent identity and access management (IAM) services to CSPs and cloud customers, often as an intermediary. This can take the form of a variety of services, including single sign-on, certificate management, and cryptographic key escrow.

Regulators The entities that ensure organizations are in compliance with the regulatory framework for which they're responsible. These can be government agencies, certification bodies, or parties to a contract. Regulations include the Health Information Portability and Accountability Act (HIPAA), the Graham-Leach-Bliley Act (GLBA), the Payment Card Industry Data Security Standards (PCI-DSS), the International Organization for Standardization (ISO), the Sarbanes–Oxley Act (SOX), and so forth. Regulators include the Federal Trade Commission (FTC), the Securities and Exchange Commission (SEC), and auditors commissioned to review compliance with contracted or asserted standards (such as PCI-DSS and ISO), among many others.

Cloud Computing Definitions

Because cloud definitions are at the heart of understanding the following chapters and applying security fundamentals for the Certified Cloud Security Professional, we have included some of those definitions here.

Apache Cloud Stack 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 Requirement An operational driver for decision making and input for risk management.

Cloud App (Cloud Application) The phrase used to describe a software application accessed via the Internet; may include an agent or applet installed locally on the user's device.

Cloud Architect Subject matter expert for cloud computing infrastructure and deployment.

Cloud Backup Backing up data to a remote, cloud-based server. As a form of cloud storage, cloud backup data is stored in an accessible form from multiple distributed resources that comprise a cloud.

Cloud Computing A type of computing, compared to grid computing, that relies on ensuring computing resources rather than having local server or personal devices to handle applications. The goal of cloud computing is to apply traditional supercomputing or high-performance computing power, normally used by military and research facilities, to perform tens of trillions of computations per second and consumer-oriented applications such as financial portfolios, or even to deliver personalized information or power immersive computer games.

Cloud Computing Reseller A company that purchases hosting services from a cloud server hosting or computing provider and then resells them to its own customers.

Cloud Migration The process of transitioning all or part of a company's data, applications, and services from onsite premises 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 The ability to move applications and associated data between one cloud provider and another, or between legacy and cloud environments.

Cloud Provider A service provider that offers customer storage or software solutions available via a public network, usually the Internet. The cloud provider dictates both the technology and operational procedures involved.

Cloud Services Broker (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 cloud service providers. It acts as a liaison between cloud services customers and cloud service providers, selecting the best provider for each customer and monitoring the services.

Cloud Storage The storage of data online in the cloud, wherein a company's data is stored in an accessible form from multiple distributed and connected resources that comprise a cloud.

Cloud Testing Load and performance testing conducted on the applications and services provided by a cloud provider, particularly the capability to access the services, in order to ensure optimal performance and scalability under a wide variety of conditions.

Community Cloud A model where the cloud infrastructure is designed for use by a specific community. Generally, this is a community of users or consumers with shared concerns, missions, and/or security requirements.

Enterprise Application The term used to describe applications or software that a business would use to assist the organization in solving enterprise problems.

Eucalyptus An open source cloud computing and Infrastructure as a Service (IaaS) platform for enabling private clouds.

FIPS 140-2 A NIST document that lists accredited and outmoded cryptosystems.

Hybrid Cloud A cloud solution that mixes elements of public, private, and community cloud models.

Infrastructure as a Service (IaaS) One of three main categories of cloud computing services, alongside Software as a Service (SaaS) and Platform as a Service (PaaS). Offers only hardware and administration, leaving the customer responsible for the OS and other software.

Managed Service Provider An IT service where the customer dictates both the technology and operational procedures, and an external party executes administration and operational support according to a contract.

Multi-Tenant Multiple customers using the same public cloud (and often the same hosts, in a virtualized cloud environment).

NIST 800-53 A guidance document with the primary goal of ensuring that appropriate security requirements and controls are applied to all U.S. federal government information in information management systems.

Platform as a Service (PaaS) A way for customers to rent hardware, operating systems, storage, and network capacity over the Internet from a cloud service provider. PaaS is one of three main categories of cloud computing services, alongside Software as a Service (SaaS) and Infrastructure as a Service (IaaS).

Private Cloud The phrase used to describe a cloud computing platform that is implemented within the organization. A private cloud is designed to offer the same features and benefits of public cloud systems but removes a number of objections to the cloud computing model, including control over enterprise or customer data, worries about security, and issues connected to regulatory compliance or contractual agreements.

Software as a Service (SaaS) SaaS is a software delivery method that provides access to software and its functionality remotely as a web-based service. Software as a Service allows organizations to access business functionality at a cost typically less than paying for licensed application because SaaS pricing is based on a monthly fee. SaaS is one of three main categories of cloud computing services, alongside Platform as a Service (PaaS) and Infrastructure as a Service (IaaS).

Trusted Cloud Initiative (TCI) Reference Model The TCI reference model is a guide for cloud providers, allowing them to create a holistic architecture (including the physical

facility of the datacenter, the logical layout of the network, and the processes necessary to utilize both) that cloud customers can purchase and use with comfort and confidence. For more information, visit `https://cloudsecurityalliance.org/wp-content/uploads/2011/10/TCI-Reference-Architecture-v1.1.pdf`.

Vendor Lock-in Vendor lock-in occurs in a situation where a customer may be unable to leave, migrate, or transfer to an alternate provider due to technical or nontechnical constraints.

Vendor Lock-out Vendor lock-out occurs when a customer is unable to recover or access their own data due to the cloud provider going into bankruptcy or otherwise leaving the market.

Virtualization Creating a virtual (a logical vs. a physical) version of something, including virtual computer hardware platforms, operating systems, storage devices, and computer network resources. Computer hardware virtualization is a way of improving overall efficiency. It involves CPUs that provide support for virtualization in hardware, and other hardware components that help improve the performance of a guest environment.

The successful CCSP candidate will be familiar with each of these terms. We will go into more detail regarding these terms over the course of the book.

Back to Basics

It's also important to remember all the security fundamentals used throughout the industry. For instance, the familiar CIA triad will be mentioned extensively throughout the CBK, the exam, and this book.

- **Confidentiality**: Protecting information from unauthorized access/dissemination
- **Integrity**: Ensuring that information is not subject to unauthorized modification
- **Availability**: Ensuring that authorized users can access the information when they are permitted to do so

Foundational Concepts of Cloud Computing

There are some aspects of cloud computing that are pervasive throughout all discussion of the topic. We're introducing them here, and you should become familiar with them. These concepts will be included in various discussions throughout the book.

Sensitive Data

Each organization will have its own risk appetite and desire for confidentiality. No matter how each cloud customer makes their own determination for these aspects of their data, the cloud provider must offer some way for the customer to categorize data according to its sensitivity, and sufficient controls to ensure these categories are protected accordingly.

Virtualization

Virtualization is one of the technologies that has made cloud services a financially viable business model. Cloud providers can purchase and deploy a sufficient number of hosts for a respective number of customers and users without wasting capacity or letting resources go idle.

In a virtualized environment, a cloud user can log onto the cloud network and boot up a synthetic version of a desktop computer. To the user, there is no appreciable difference between the virtual machine (VM) and a traditional computer. However, from the provider's perspective, the VM being offered to the user is just a piece of software, not an actual, dedicated piece of hardware being exclusively operated by the user. Indeed, there may be several, or even dozens, of VMs operating on a single host in the cloud space concurrently. When the user logs off or shuts down, the cloud network takes a snapshot of the user's VM, capturing it as a single file that can be stored somewhere else in the cloud until the user next requests access, when the VM can be restored exactly as they left it.

In this way, the cloud provider can offer services to any number of customers and users, and not be required to purchase a new hardware device for each new user. This economy of scale allows the cloud provider to offer the same basic IT services that the users expect from traditional networks with much less cost and at an enhanced level of service.

There are many virtualization product vendors, including VMware and Microsoft. There are also a variety of implementation strategies, and two fundamental virtualization types (Type 1 and Type 2).

Encryption

As an IT security professional, you should already be familiar with the basic concepts and tools of encryption. However, in terms of cloud services, encryption plays an enhanced role and presents some additional challenges.

Because your cloud data will be in an environment controlled and operated by personnel other than your organization, encryption offers a degree of assurance that nobody without authorization will be able to access your data in a meaningful way. You can encrypt your data before it reaches the cloud, and only decrypt it as necessary.

Another concern related to cloud operation is that it necessitates remote access. As with any remote access, there will always be a risk (however great or slight) of interception of data, eavesdropping, and man-in-the-middle attacks. Encryption also assists in alleviating this concern by mitigating this threat to some degree; if data in motion is encrypted, it is that much more difficult to access even if it is intercepted.

Auditing and Compliance

Cloud services pose specific challenges and opportunities for regulatory compliance and auditing.

From a compliance perspective, service providers may be able to offer holistic solutions for organizations under particular regulatory schema. For instance, the cloud provider may have an extant, known, tested control set and procedural outline for PCI, HIPAA, or GLBA. This could be extremely appealing to potential customers, as the difficulty and effort expended in trying to stay compliant can now be shifted out of the customer organization and over to the provider.

Conversely, auditing becomes more difficult. Cloud providers are extremely reluctant to allow physical access to their facilities or to share network diagrams and lists of controls; maintaining confidentiality of these things enhances the provider's overall security. However, these are essential elements of an audit. Also, as you'll see in upcoming chapters, it is difficult to determine exactly where in the cloud environment a given organization's data is located at any moment, or which devices contain a certain customer's data, making auditing even more difficult. Audits will require the cooperation of the cloud provider, and providers have thus far disallowed the requisite level of access for the purpose. Instead, cloud providers often offer an assertion of their own audit success (often in the form of a Statement on Standards for Attestation Engagements and Service Organization Controls Type 3 report). Any organization considering cloud migration should confer with the regulatory agencies that provide oversight for them in order to determine whether this limited audit insight will be sufficient for the regulators.

Cloud Service Provider Contracts

The business arrangement between the cloud provider and the cloud customer will usually take the form of a contract and a service-level agreement (SLA). The contract will spell out all the terms of the agreement: what each party is responsible for, what form the services will take, how issues will be resolved, and so on. The SLA will set specific, quantified goals for these services and their provision over a certain timeframe.

For instance, the contract might stipulate "The Provider will ensure the Customer has constant, uninterrupted access to the Customer's data storage resources." The SLA will then explicitly define the metrics for what "constant, uninterrupted access" will mean: "There will be no interruption of connectivity to data storage longer than three (3) seconds per calendar month." The contract will also state what the penalties are (usually financial) when the cloud provider fails to meet the SLA for a given period: "Customer's monthly fee will be waived for any period following a calendar month in which any service level has not been attained by Provider."

These are obviously rough examples, but they demonstrate the relationship between the contract, the SLA, the cloud provider, and the cloud customer. The book will continually refer to the contract and the SLA based on the relationship explained here.

Summary

In this chapter, we have explored business requirements, cloud definitions, cloud computing roles and responsibilities, and foundational concepts of cloud computing. As this is the introductory chapter, we will explore each of these topics in more detail as we move ahead.

Exam Essentials

Understand business requirements. Always bear in mind that all management decisions are driven by business needs, including security and risk decisions. Security and risk should be considered before these decisions are made, and may not take precedence over the business and operational requirements of the organization.

Understand cloud vernacular and definitions. Make sure you have a clear understanding of the definitions introduced in Chapter 1. A great deal of the CCSP exam focuses on terms and definitions.

Be able to describe the cloud service models. It is vitally important that you understand the differences between the three cloud service models, IaaS, PaaS, and SaaS, and the different features associated with each.

Understand cloud deployment models. It is also important for you to understand the features of each of the four cloud deployment models, Public, Private, Community, and Hybrid, as well as their differences.

Be familiar with cloud computing roles and the associated responsibilities. Make sure you know and understand the different roles and the responsibilities of each of the roles. We will explore roles in more detail in the chapters that follow.

Written Labs

1. Go to the CSA website and watch the video titled "Intro to Cloud Computing" at `https://cloudsecurityalliance.org/education/white-papers-and-educational-material/courseware/`. When you are done, spend some time exploring the site.
2. Write down three things you can think of that might be legitimate business drivers for an organization considering cloud migration.
3. List the three cloud computing service models and the advantages and disadvantages of each.

Review Questions

You can find the answers in Appendix A.

1. Which of the following is *not* a common cloud service model?
   - **A.** Software as a Service
   - **B.** Programming as a Service
   - **C.** Infrastructure as a Service
   - **D.** Platform as a Service
2. All of these technologies have made cloud service viable except:
   - **A.** Virtualization
   - **B.** Widely available broadband
   - **C.** Cryptographic connectivity
   - **D.** Smart hubs
3. Cloud vendors are held to contractual obligations with specified metrics by:
   - **A.** 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. Cloud Access Security Brokers (CASBs) might offer all the following services EXCEPT:
   - **A.** Single sign-on
   - **B.** BC/DR/COOP
   - **C.** IAM
   - **D.** Key escrow

7. Encryption can be used in various aspects of cloud computing, including all of these except:
   - **A.** Storage
   - **B.** Remote access
   - **C.** Secure sessions
   - **D.** Magnetic swipe cards

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 services
   - **B.** Negating the need for backups
   - **C.** Resource pooling
   - **D.** Measured or metered service

10. All of the following can result in vendor lock-in except:
    - **A.** Unfavorable contract
    - **B.** Statutory compliance
    - **C.** Proprietary data formats
    - **D.** Insufficient bandwidth

11. The risk that a cloud provider might go out of business and the cloud customer might not be able to recover data is known as:
    - **A.** Vendor closure
    - **B.** Vendor lock-out
    - **C.** Vendor lock-in
    - **D.** Vending route

12. All of these are features of cloud computing except:
    - **A.** Broad network access
    - **B.** Reversed charging configuration
    - **C.** Rapid scaling
    - **D.** On-demand self-service

13. When a cloud customer uploads PII to a cloud provider, who becomes 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.** BIA
    - **C.** RMF
    - **D.** CIA triad

15. The cloud deployment model that features organizational ownership of the hardware and infrastructure, and usage only by members of that organization, is known as:
    - **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. If a cloud customer wants a secure, isolated sandbox in order to conduct software development and testing, which cloud service model would probably be best?
    - **A.** IaaS
    - **B.** PaaS
    - **C.** SaaS
    - **D.** Hybrid

19. If a cloud customer wants a fully-operational environment with very little maintenance or administration necessary, which cloud service model would probably be best?

A. IaaS

B. PaaS

C. SaaS

D. Hybrid

20. If a cloud customer wants a bare-bones environment in which to replicate their own enterprise for BC/DR purposes, which cloud service model would probably be best?

A. IaaS

B. PaaS

C. SaaS

D. Hybrid