Chapter 16

Cloud Computing

Network+ Certification Exam Objectives

The Network+ certification exam has expanded coverage of cloud computing substantially with the latest version of the exam. That is why this book now has an entire chapter devoted to cloud.

1.8: Summarize cloud concepts and connectivity options
Deployment Models

Learning Outcomes

16.1 Recall cloud concepts.
16.2 Explain multitenancy and related issues.
16.3 Discuss various cloud deployment models.
16.4 Understand service models.
16.5 Summarize major cloud security issues.
16.6 Provide examples of specific cloud implementations.

Key Terms

Artificial Intelligence as a Service (AIaaS)
audit monitor
cloud computing
Content as a Service (CaaS)
Data as a Service (DaaS)
Desktop as a Service
elasticity
Function as a Service (FaaS)
high-performance cloud (HPC)
hypervisor
Infrastructure as a Service (IaaS)
logical network perimeter
Mobile Backend as a Service (MBaaS)
multicloud
multitenancy
Platform as a Service (PaaS)
polycloud
Security as a Service (SECaaS or SaaS)
Software as a Service (SaaS)
scalability
virtual desktop environment (VDE)
virtual desktop infrastructure (VDI)
virtual storage

Overview

This chapter explores cloud computing, beginning with the concept of virtualization and expanding that to the modern cloud landscape. The various deployment and service models are thoroughly explored, as well as connectivity options. Topics such as multitenancy and elasticity are discussed. Finally, security implications are explored.

16.1 Cloud Computing Concepts

1.8 NET

Cloud computing is a general term used to describe a new class of network-based computing that takes place over the Internet, a collection or group of integrated and networked hardware, software, and Internet infrastructure (called a platform). Using the Internet for communication and transport provides hardware, software, and networking services to clients. Cloud computing is often preferred because of its ability to scale up or down quickly, referred to as **elasticity**.

In all clouds, someone else is providing the physical machines. Users concerned about power, bandwidth, maintenance, physical security, or (sometimes) scaling can find cloud computing to be the solution. And just as important, users only pay for what they use. Cloud computing uses servers distributed geographically. In some cases, the servers are in other countries. This brings the benefit of fault tolerance, but has some security concerns:

- Privacy laws vary by region.
- Ensuring that a customer's data is segregated from other customers' data is the primary data protection issue.

The National Institutes of Standards and Technology defines a cloud as follows:

> 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. This cloud model is composed of five essential characteristics, three primary service models, and four deployment models. (NIST Special Publication 800-145)

The five characteristics are:

- *On-demand self-service.* This means customers can provision resources as needed, which allows them a tremendous amount of flexibility. In fact, on-demand self-service is one of the major driving forces behind the growth of cloud computing.
- *Broad network access.* Cloud resources are available over the network and accessed through a standard network. This makes it easy to access the resources from anywhere in the world at any time.
- *Resource pooling.* The provider's computing resources are pooled to serve multiple consumers. This pooling is accomplished via virtualization, which will be discussed in detail in the next section. Resource pooling also involves multitenancy. **Multitenancy** simply means that multiple customers are using the same resources. Multitenancy is described in more detail later in this section.
- *Rapid elasticity.* The resources can be scaled up or down in a very short period. This is sometimes called **scalability**. This allows customers to pay for more services only when and if they need them. That is sometimes called *pay as you grow*.
- *Measured service.* Cloud systems automatically control and optimize resource use. This is usually done via metering. One customer is not permitted to monopolize resources.

This description of the five essential features is fairly standard across the cloud-computing industry. Even for those who don't reside in the United States, the US National Institute of Standards is a good resource. The service models and deployment models will be described later in this chapter.

As early as 1993, the term *cloud* was used in reference to distributed computing. However, cloud computing did not get really started until well into the 2000s. In 2006, Amazon released its Elastic Compute Cloud (EC2) specifically for Infrastructure as a Service (IaaS). In 2010, Microsoft launched their cloud solution, Azure. Then in 2011, IBM announced its IBM Smartcloud. In 2012, Oracle launched the Oracle Cloud. Cloud services have continued to grow since that time.

NET 1.8

16.2 Multitenancy

Multitenancy is listed as part of the NIST 5 characteristics of cloud computing, and it is a critical topic in cloud computing. When you have cloud architecture, there are groups of users, typically called tenants, that have common access privileges. This is different than when there are completely separate software instances. In multitenancy, several customers have the same application running in the same operating system. This approach provides a substantial cost savings. It can also be beneficial for activities like data mining. With the aggregation of many users data, data mining is more effective. Multitenancy also simplified release management, since all users share the same application.

Introduction to Virtualization

The concept of virtualization includes not only cloud solutions, but local virtual machines. One might have a virtual machine such as Oracle Virtual Box or VMWare workstation on one's desktop. Inside that virtual machine software one can host other operating systems. The primary difference between the local virtual machine and the cloud solution is the source of the underlying hardware. With a local computer virtual machine, a subset of the computer's resources (hard drive storage, RAM, etc.) is segmented to create the virtual machine. With a cloud solution, many servers' resources (again, hard drive storage, RAM, etc.) are all pooled into a single virtualized pool and then assigned to individual virtual machines. In both situations there are some common technologies:

- *Virtual storage.* The virtual servers are hosted on one or more physical servers. The hard drive space and RAM of those physical servers are partitioned for the various virtual servers' usage, providing **virtual storage**.
- *Audit monitor.* There is usually an **audit monitor** that monitors usage of the resource pool. This monitor will also ensure that one virtual server does not/cannot access data of another virtual server.
- *Hypervisor.* The **hypervisor** mechanism is the process that provides the virtual servers with access to resources. This is also called a virtual machine monitor (VMM).
- *Logical network perimeter.* Since the cloud consists of virtual servers, not physical ones, there is a need for a logical network and a **logical network perimeter**. This perimeter isolates resource pools from each other.

It is important to remember that these components must exist in any virtualized environment, whether that environment is on a single computer or a cloud. Virtualization is the backbone of cloud technology. It is how a diverse set of hardware, perhaps distributed geographically, can appear to be a monolithic resource pool. Then that monolithic resource pool, in turn, appears to be individual hardware to the customer. The relationship between a customer's computer and the data centers in a cloud is shown in Figure 16-1.

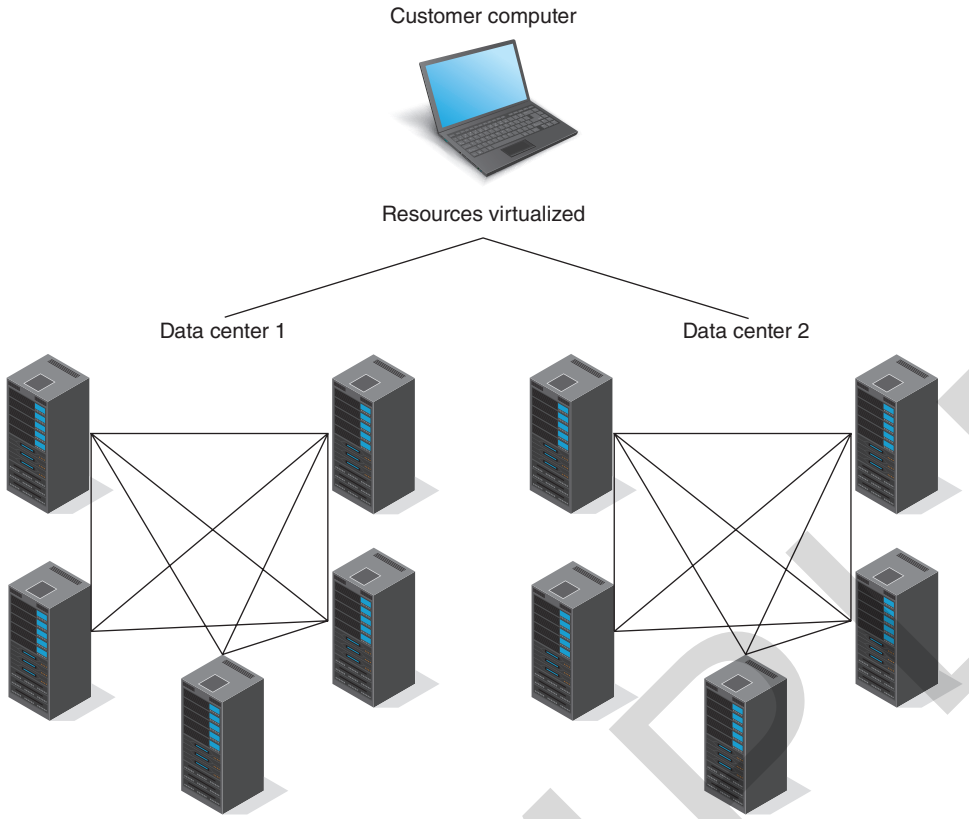

Goodheart-Willcox Publisher; (workstation icon) romvo/Shutterstock.com; (server icons) aShatilov/Shutterstock.com

Figure 16-1 The relationship between the customer and cloud resources that are virtualized.

Related to cloud computing is fog computing. Fog computing is using decentralized servers in between network core and network edge for data processing. It is also frequently used with IoT, particularly in corporate and industrial settings.

16.3 Deployment Models

1.8 NET

While the technical details of cloud computing are important, so are the deployment models. This concerns who owns the cloud, who is responsible for the administration of the cloud, and who can access the cloud.

Public clouds are defined by the NIST as simply clouds that offer their infrastructure or services to either the general public or at least a large industry group. These are the clouds most people are familiar with. Microsoft Azure, Amazon Web Service (AWS), and Google Cloud are well-known examples. Any company, charitable organization, or even an individual can sign up to use these cloud provider services. Then the customer pays only for the resources used.

Private clouds are those clouds used specifically by a single organization without offering the services to an outside party. Hewlett Packard Enterprise (HPE) is a well-known private cloud provider. VMWare also offers private cloud solutions. Some organizations completely host and manage their own private cloud. That does require substantial resources. There are, of course, hybrid clouds, which combine the elements of a private and public cloud. These are essentially private clouds that have some limited public access.

As the name suggests, a hybrid cloud is a composition of a public cloud with a private environment. Google defines a hybrid cloud as "... one in which applications

are running in a combination of different environments. Hybrid cloud computing approaches are widespread because almost no one today relies entirely on the public cloud."

Community clouds are a midway point between private and public. These are systems wherein several organizations share a cloud for specific community needs. For example, several computer companies might join to create a cloud devoted to common security issues.

The public, private, hybrid, and community models are the ones most focused on by the NIST and other sources. In addition to the deployment models, there are models regarding how one uses cloud services. **Multicloud** is a common use case. Multiple cloud vendors are used heterogeneously. This mitigates dependency on a single vendor. Cloud assets (applications, virtual servers, etc.) are hosted across multiple public clouds. One can also include private clouds in the architecture. The issue in multicloud is not whether you are using a public or private cloud. Rather, the issue is whether or not your resources are hosted by a single cloud. Cloud computing itself is quite resilient. Using a multicloud approach further increases your stability and reliability.

Polycloud is similar multicloud, but in this case the different public clouds are being utilized not for flexibility and redundancy, but rather for specific services each provider offers. For example, you may wish to use Microsoft Azure for its business analytics offerings, while using Oracle Cloud for database services. While the large cloud providers tend to offer a wide range of services, you will often find that one particular provider is better at a particular application than another.

Another application of cloud computing is **high-performance cloud (HPC)**. HPC is the use of cloud services for high-performance computing. Such HPC applications would normally require clusters of computers or a supercomputer. There are several companies, including Amazon Web Services, Google, and Microsoft, that offer access to HPC.

NET 1.8

16.4 Service Models

There are numerous service models, and it seems like more are appearing each day. The NIST only defines three: **Software as a Service (SaaS)**, **Platform as a Service (PaaS)**, and **Infrastructure as a Service (IaaS)**. We will examine these three first, but will also examine additional service models.

The NIST definition of SaaS is as follows:

> The capability provided to the consumer is to use the provider's applications running on a cloud infrastructure. The applications are accessible from various client devices through either a thin client interface, such as a web browser (e.g., web-based email), or a program interface. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, storage, or even individual application capabilities, with the possible exception of limited user-specific application configuration settings.

The NIST definition of PaaS is as follows:

> The capability provided to the consumer is to deploy onto the cloud infrastructure consumer-created or acquired applications created using programming languages, libraries, services, and tools supported by the provider. The consumer does not manage or control the underlying cloud infrastructure including network, servers, operating systems, or storage, but has control over the deployed applications and possibly configuration settings for the application-hosting environment.

NIST defines IaaS as:

> ...where the consumer is able to deploy and run arbitrary software, which can include operating systems and applications. The consumer does not manage or control the underlying cloud infrastructure but has control over operating systems, storage, and deployed applications; and possibly limited control of select networking components (e.g., host firewalls).

There are numerous variations of the *as a service* model. Each of these provides some particular service via cloud computing. These are becoming more important for modern organizations. Most organizations use some service, and if it is a computer related service, it is likely they are getting it via a cloud. Several of these are discussed in the following paragraphs:

Artificial Intelligence as a Service (AIaaS), as the name suggests, allows one to access machine learning and artificial intelligence resources via cloud computing. These services can provide speech to text, data classification, bots, and other AI/ML resources.

Data as a Service (DaaS) could be thought of as a specialized IaaS. The consumer does not need the entire infrastructure provided via the cloud, but rather the data source. This can be a relational database or a NoSQL datastore or, in fact, any type of data source. The key point is that the data as well as data management functions are all hosted by a cloud provider.

Desktop as a Service is a common service model. This goes by various names, such as **virtual desktop infrastructure (VDI)**. As the name suggests, DaaS hosts a desktop operating system on a centralized server in a data center. There are two types of DaaS. A persistent DaaS provides each user with their own desktop image, which can be customized and saved for future use. With a nonpersistent DaaS, there exists a pool of uniform desktops that users can access when needed. VDI is sometimes referred to as **VDE** or **virtual desktop environment**.

Mobile Backend as a Service (MBaaS) allows mobile apps to be linked to cloud storage and resources. This is used for push notifications, user management, and other services.

Security as a Service (SECaaS or SaaS) provides cybersecurity services. As security needs become more complex, there is often a need to outsource some or all of the security to a specialized provider. SECaaS allows the customer to access only the security services needed. One might use SECaaS for threat intelligence or antimalware, or frankly any cybersecurity function.

Another new concept has been **Function as a Service (FaaS)**. This also sometimes called serverless computing. While there clearly are servers involved at some point in the process, from the customer perspective there aren't. The customer does not purchase a server, not even a virtual one. Rather the customer rents specific services.

Content as a Service (CaaS), as the name suggests, provides content. This can be done as a web service. Often, the users of CaaS integrate such content into their own applications and/or websites. Developers can access content usually via an application programming interface (API) or web service, and then integrate that content as needed.

Regardless of the model, one must interact with the cloud resources. There are three primary ways to do this. Perhaps the most common is a graphical user interface (GUI). This can be through an application on a computer, or often a website interface. The second method is a command line interface (CLI) using a Linux shell or Windows command line window. The third is an API. This option allows the consumer to program their own interface using the API.

NET 1.8

16.5 Security Issues

It is undeniable that cloud computing provides numerous benefits. Those benefits include scalability, affordability, and efficiency. However, cloud computing also presents security issues. In some cases, there are security vulnerabilities in the underlying cloud software. In early 2022 there was a vulnerability labeled CVE-2022 that allowed remote code execution in Spring Cloud functions.

As more and more organizations embrace cloud computing, this also leads to cloud security issues. In this section we will review cloud security issues, standards for cloud security, as well as general security guidance. Any distributed system, such as a cloud, has certain technical challenges. These include:

- Synchronization: multiple clocks (difficult to agree on exact time)
- Concurrency: multiple simultaneous accesses potentially conflicting
- Failures: high probability of failures for individual components
- Consensus: difficult to reach consensus with shared data storage

Concurrency and consensus are related issues. Consider four different data centers, in four different cities. This is shown in Figure 16-2.

The customer simply wishes to access their data. As their employees travel, they will most likely be connected to the nearest data center, but it is important that they access the same data. An employee connecting to the cloud in California should not see different data than an employee connecting to the cloud in Asia. It is important that data be replicated between data centers to ensure that the data is the same in all data centers. This is also one of the disaster recovery benefits of using cloud computing. With data and resources in multiple geographic locations, it is difficult to imagine a disaster that would disable all those locations.

IBM defines cloud security as "... a collection of procedures and technology designed to address external and internal threats to business security. Organizations need cloud security as they move toward their digital transformation strategy and incorporate cloud-based tools and services as part of their infrastructure."

In addition to these management issues, there are of course issues with hackers, malware, security breaches, and the like. In fact, all of the security issues faced in traditional computing are also present in a cloud. A 2021 AWS survey revealed the following:

- 95 % of cybersecurity professionals confirm they are extremely to moderately concerned about public cloud security.
- The number one concern remains misconfiguration of the cloud platform (71%). Exfiltration of sensitive data came in second (59%). Insecure APIs (54%) rounded out the top three concerns.
- 58% use periodic vulnerability and compliance reports as the primary method of communication with system owners about security and compliance issues needing remediation. This is followed by automatically opened tickets (47%) using tools such as Jira, ServiceNow, etc.
- >40% of organizations embrace hybrid cloud (44%) and multicloud deployments (43%) for planned redundancy because of commitments to legacy applications in traditional data centers. Single cloud deployments (11%) continue to diminish in importance.
- 90% of organizations use more than two cloud providers.
- 66% of organizations prioritize cost-effectiveness as a leading criterion when selecting a cloud security provider, followed by scalability (52%), ease of deployment (51%), and tools that can be deployed with automation (48%). (2021 AWS Cloud Security Report)

While cloud computing still poses the same security issues as traditional on-premise networks, there are some issues of security that change in importance

Figure 16-2 Geographically distributed data centers.

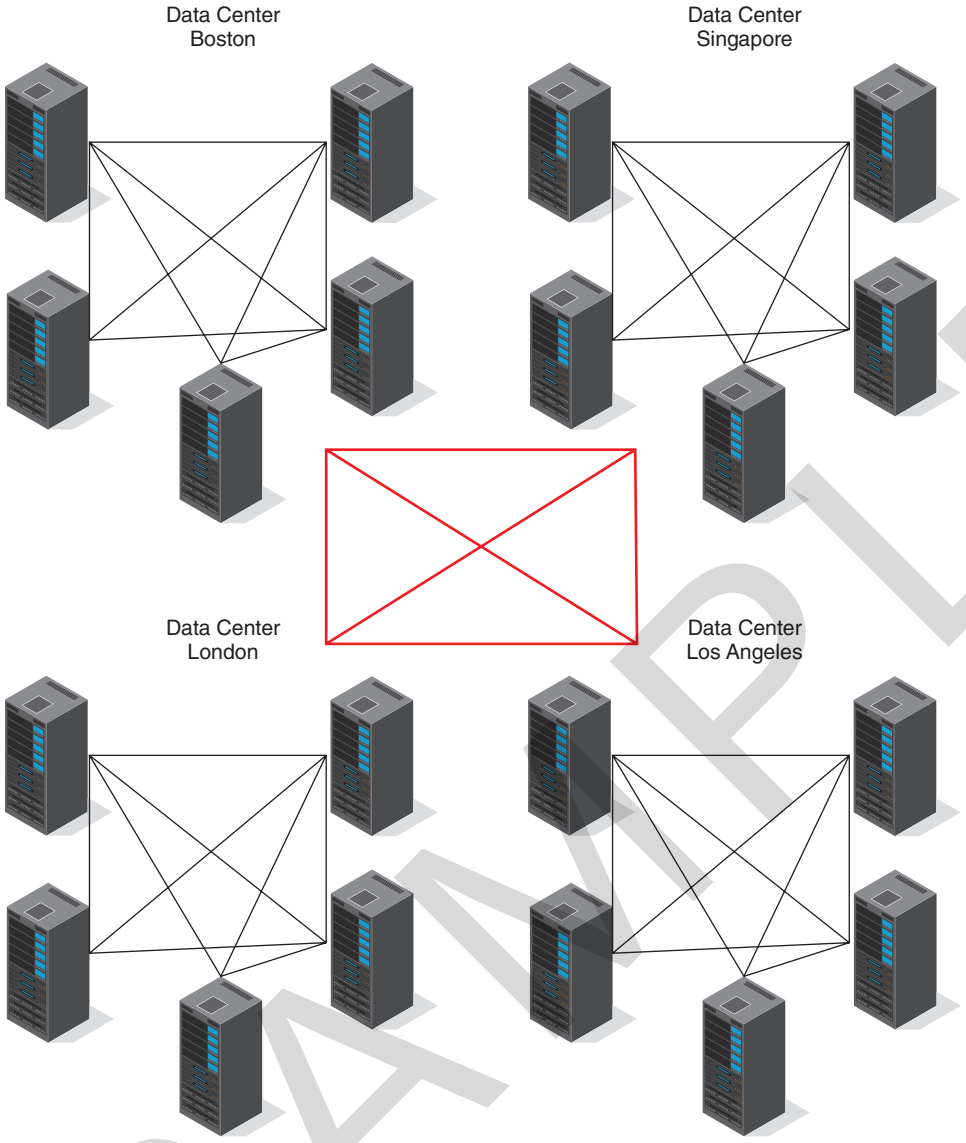

Goodheart-Willcox Publisher; (server icons) aShatilov/Shutterstock.com

with respect to cloud computing. The first issue is that data from multiple customers is stored in the same cloud, unless you are using a private cloud. This is the previously mentioned issue of multitenancy. Cloud providers generally handle this for you, but you should inquire as to how they handle this issue before committing to a cloud vendor.

Availability is also a concern. One of the reasons cloud computing is so popular is that it should increase availability. You should inquire with any cloud vendor details on how they maintain availability. This will include questions as to where data centers are located. That, of course, leads to the issue of synchronization. When data is distributed geographically, synchronization is a nontrivial problem.

Considering cloud security may seem daunting. Fortunately, there are standards for cloud security. These can provide a starting place for you to explore cloud security.

NIST Special Publication 500-29

The NIST Special Publication 500-29 standard provides a roadmap to cloud security. It begins by covering cloud technology and cloud computing use cases. Then it delves into cloud computing standards. These are not necessarily security standards, but instead, just cloud computing standards. However, NIST 500-29 also covers cloud security including authentication, identity management, security controls, and security policy.

ISO 27017

ISO 27017 provides guidance for cloud security. It does apply the guidance of ISO 27002 to the cloud, but then adds seven new controls, listed below.

CLD.6.3.1: Agreement on shared or divided security responsibilities between the customer and cloud provider.

CLD.8.1.5: Addresses how assets are returned or removed from the cloud when the contract is terminated.

CLD.9.5.1: The cloud provider must separate the customer's virtual environment from other customers or outside parties.

CLD.9.5.2: The customer and the cloud provider both must ensure the virtual machines are hardened.

CLD.12.1.5: It is solely the customer's responsibility to define and manage administrative operations.

CLD.12.4.5: The cloud provider's capabilities must enable the customer to monitor their own cloud environment.

CLD.13.1.4: The virtual network environment must be configured so that it at least meets the security policies of the physical environment.

The following cloud providers are compliant with ISO 27017 (note: there may be more; these are simply major cloud providers that advertise their compliance):

- IBM Cloud
- Microsoft Azure
- Cisco Webex
- Google Cloud Platform
- Amazon Web Services
- Nutanix Xi Cloud

ISO 27018

ISO 27018 is closely related to ISO 27017. ISO 27018 defines privacy requirements in a cloud environment, particularly how the customer and cloud provider must protect personally identifiable information (PII). There was an original version of this standard in 2014, and it was updated in 2019.

FedRAMP

FedRAMP (Federal Risk and Authorization Management Program) involves back-end security protection for cloud services and cloud-based customer data. This standard was released in 2011 with the goal of providing guidance to US federal agencies on cloud security.

SOC

System and Organization Controls (SOC), which include SOC 1 and SOC 2, are actually financial audit/due diligence standards. These may seem removed from cloud computing but are not. Consider how much data organizations store in the cloud. At least some familiarity with SOC can be beneficial. SOC 2 reports are more related to technology. An SOC 2 report will address what are known as the trust service principles:

- The security of a service organization's system
- The availability of a service organization's system
- The processing integrity of a service organization's system
- The confidentiality of the information that the service organization's system processes or maintains for user entities
- The privacy of personal information that the service organization collects, uses, retains, discloses, and disposes of for user entities

Other Sources

The Cloud Security Alliance (CSA) is a very good source for information on cloud security. The cloud security alliance is an organization that focuses on security for cloud resources. You can find their website at https://cloudsecurityalliance.org. The website offers a number of resources, including numerous publications. They also sponsor a cloud security certificate.

There are also now cloud-specific security certifications. After you have passed the CompTIA Network+, you may wish to consider CompTIA Cloud+. That certification covers general cloud technology as well as cloud security. ISC2 also has the Certified Cloud Security Specialist certification, which is focused solely on cloud security.

16.6 Specific Cloud Implementations

The Network+ won't ask you about specific cloud offerings from particular vendors. However, it is useful for you to be familiar with the major vendors.

Oracle

Oracle has Oracle Cloud Infrastructure (OCI) to offer IaaS. There are virtual machines, load balancing, storage, edge services, and integration with existing VMWare data centers. Oracle also has Oracle Cloud Platform (OCP). This is their Platform as a Service offering. It includes application development and business analytics. Oracle also offers Oracle Cloud Applications (OCA) as their Software as a Service offering. This includes IoT applications, Enterprise Resource Planning, Analytics, and Block-Chain Cloud service.

Microsoft

Azure is Microsoft's cloud solution. Azure is comprised of 200+ physical data centers arranged into regions. These include locations in Hong Kong, Singapore, Iowa, Texas, Toronto, Ireland, Paris, Chennai, Mumbai, Doha, Madrid, and Johannesburg, among other locations. Together, these data centers form the Azure global infrastructure. (See https://azure.microsoft.com/en-us/global-infrastructure/ for more information.) A tenant in Azure is a complete entity. It is the top-level entity/object. Microsoft defines a tenant as follows:

> A tenant represents an organization. It's a dedicated instance of Azure AD that an organization or app developer receives at the beginning of a relationship with Microsoft. That relationship could start with signing up for Azure, Microsoft Intune, or Microsoft 365, for example.

> Each Azure AD tenant is distinct and separate from other Azure AD tenants. It has its own representation of work and school identities, consumer identities (if it's an Azure AD B2C tenant), and app registrations. An app registration inside your tenant can allow authentications only from accounts within your tenant or all tenants. (https://docs.microsoft.com/en-us/azure/active-directory/develop/quickstart-create-new-tenant)

Amazon

Amazon Web Services (AWS) is Amazon's cloud offering. The actual data is hosted in server farms distributed around the world. AWS currently offers about 200 different services including Amazon Aurora (a relational database), Amazon Augmented AI, Amazon DynamoDB (NoSQL Datastore), Amazon GuardDuty (threat detection), and more. Amazon, like most cloud providers, offers security-specific recommendations on their website. In fact, they have a section devoted to cloud security at https://aws.amazon.com/security/. Amazon also advertises its compliance with ISO 27017 at https://aws.amazon.com/compliance/iso-27017-faqs.

IBM Cloud

IBM advertises a number of cloud solutions including hybrid cloud offerings, AI and automation via cloud, and Infrastructure as a Service. The current IBM cloud is a rebranding of various services that IBM acquired. The rebranding occurred in 2017. As of this writing, IBM cloud offers over 170 separate services, not just the sample listed at the beginning of this paragraph. See www.ibm.com/cloud/solutions for more details.

Summary

16.1 Cloud Computing Concepts

- The use of cloud computing is growing.
- Elasticity is a primary advantage of cloud computing.
- The five essential characteristics of the cloud model are on-demand self-service, broad network access, resource pooling, rapid elasticity, and measured service.

16.2 Multitenancy

- Multitenancy is a critical topic in cloud computing.
- In multitenancy, several customers have the same application running in the same operating system.
- Virtualization is the fundamental technology used in cloud computing.

16.3 Deployment Models

- Deployment models concern who owns the cloud, who is responsible for administration of the cloud, and who can access the cloud.
- Major deployment models include public, private, hybrid, and community clouds.

16.4 Service Models

- There are multiple service models including PaaS, IaaS, DaaS, SECaaS, MBaaS, and others.
- New service models are constantly being added.

16.5 Security Issues

- The cloud has certain technical challenges and vulnerabilities that make security a highly important issue.
- Cloud security standards include ISO 27017, ISO 27018, and NIST SP 500-29.

16.6 Specific Cloud Implementations

- Major cloud vendors include Oracle, Microsoft, Amazon, and IBM Cloud.

Review Questions

1. Explain the concept of elasticity as it relates to cloud computing.
2. What are the five characteristics of cloud computing?
3. Define multitenancy.
4. Describe the four common technologies used in virtualization.
5. What are the four deployment models most focused on by the NIST?
6. What is the difference between polycloud and multicloud?
7. Explain the three service models defined by the NIST.
8. What are the two types of Desktop as a Service?
9. What are some common technical challenges of cloud computing?
10. What does ISO 27018 focus on?
11. List the major vendors that provide cloud services.

Sample Network+ Exam Questions

1. Janine has been tasked with identifying different cloud vendors for her company to use. The goal is to achieve flexibility and redundancy. What cloud solution is best for this purpose?
   A. Polycloud
   B. Community cloud
   C. Multicloud
   D. Public cloud
2. _____ defines privacy requirements in a cloud environment, particularly how the customer and cloud provider must protect personally identifiable information (PII).
   A. ISO 27018
   B. ISO 27017
   C. ISO 14117
   D. ISO 14708-1
3. ISO 27017 is a standard for cloud security. It adds seven new controls to ISO 27002. Which of the following is *not* one of those seven?
   A. It is the customer's responsibility to define and manage administrative operations.
   B. Cloud providers are responsible for strong cryptography being implemented.
   C. Cloud providers must enable customers to monitor their own cloud environment.
   D. There must be agreement on shared or divided security responsibilities.
4. Juan is an IT administrator for a midsize bank. His company is looking for a cloud solution. The requirements are to find particular services, even if they come from different cloud vendors. Redundancy is not a concern. What type of cloud should Juan consider?
   A. Multicloud
   B. Public cloud
   C. Hybrid cloud
   D. Polycloud

5. Mohanned is an IT manager. He is evaluating different cloud providers and is concerned about issues that are common to all distributed systems. Which of the following is *not* a common distributed system issue?
   A. Synchronization
   B. Concurrency
   C. Consensus
   D. Availability
6. What cloud model delivers server hardware with no operating system?
   A. PaaS
   B. IaaS
   C. SaaS
   D. DaaS
7. Which cloud characteristic allows you to pay for only the services used?
   A. Pay as you grow
   B. Metering
   C. Chargeback
   D. Bursting
8. What cloud model delivers all services except the application?
   A. DaaS
   B. PaaS
   C. IaaS
   D. SaaS
9. What are common management interfaces that are used to migrate and manage cloud-based resources? (Choose two.)
   A. CLI
   B. SSH
   C. GUI
   D. SNMP
10. What is a report that is most related to the processing integrity of an organization's system?
    A. SOC 1
    B. ISO 27001
    C. FIPS 140-2
    D. SOC 2

Laboratory Activity

iCanCloud Simulator

After completing this laboratory activity, you will be able to:

- Utilize the iCanCloud simulator.
- Understand cloud management.
- Explain cloud configuration.

Introduction

iCanCloud is a cloud simulator that models and simulates cloud computing systems. Go to the iCanCloud home page at https://www.arcos.inf.uc3m.es/old/icancloud/Home.html and review its features. The iCanCloud download comes with a GUI as shown below. iCanCloud and iCanCloudGUI can be downloaded from https://www.arcos.inf.uc3m.es/old/icancloud/downloads.html.

Goodheart-Willcox Publisher

The GUI is a .Jar file. If you are not familiar with .Jar files, these are Java applications. You can get more details on .Jar files at https://www.softwaretestinghelp.com/how-to-open-a-jar-file/#How_To_Open_JAR_Files.

You expand any category on the left and right click to edit that item. This is shown in the screenshot below.

Goodheart-Willcox Publisher

Equipment and Materials

- Computer with Internet access (A server is not required for this laboratory activity.)

Procedure

1. Report to your assigned workstation.
2. Boot the computer and verify it is in working order.
3. Using the documentation, create a cloud simulation with a large-scale data set.
4. Edit two clouds, and check all options.
5. Edit at least one node.
6. Add at least one connection to at least one data center.
7. When finished, close the simulation and leave your computer in the condition as specified by your instructor.
8. Answer the review questions.

Laboratory Review Questions

1. What are some of the features of iCanCloud?
2. What is the function of the iCanCloudGUI?
3. What are .Jar files, and how can they be opened?