The Official CompTIA Cloud+ Study Guide (Exam CV0-003)

Acknowledgments



Damon Garn, Author Thomas Reilly, Senior Vice President, Learning Katie Hoenicke, Senior Director, Product Management Evan Burns, Senior Manager, Learning Technology Operations and Implementation James Chesterfield, Manager, Learning Content and Design Becky Mann, Director, Product Development Katherine Keyes, Content Specialist

Notices

Disclaimer

While CompTIA, Inc., takes care to ensure the accuracy and quality of these materials, we cannot guarantee their accuracy, and all materials are provided without any warranty whatsoever, including, but not limited to, the implied warranties of merchantability or fitness for a particular purpose. The use of screenshots, photographs of another entity's products, or another entity's product name or service in this book is for editorial purposes only. No such use should be construed to imply sponsorship or endorsement of the book by nor any affiliation of such entity with CompTIA. This courseware may contain links to sites on the Internet that are owned and operated by third parties (the "External Sites"). CompTIA is not responsible for the availability of, or the content located on or through, any External Site. Please contact CompTIA if you have any concerns regarding such links or External Sites.

Trademark Notice

CompTIA®, Cloud+®, and the CompTIA logo are registered trademarks of CompTIA, Inc., in the U.S. and other countries. All other product and service names used may be common law or registered trademarks of their respective proprietors.

Copyright Notice

Copyright © 2021 CompTIA, Inc. All rights reserved. Screenshots used for illustrative purposes are the property of the software proprietor. Except as permitted under the Copyright Act of 1976, no part of this publication may be reproduced or distributed in any form or by any means, or stored in a database or retrieval system, without the prior written permission of CompTIA, 3500 Lacey Road, Suite 100, Downers Grove, IL 60515-5439.

This book conveys no rights in the software or other products about which it was written; all use or licensing of such software or other products is the responsibility of the user according to terms and conditions of the owner. If you believe that this book, related materials, or any other CompTIA materials are being reproduced or transmitted without permission, please call 1-866-835-8020 or visit https://help.comptia.org.

Table of Contents

Lesson 1: Understanding Cloud Concepts1
Topic 1A: Recognize Cloud Concepts2
Topic 1B: Recognize Cloud Terms18
Topic 1C: Understand the Troubleshooting Methodology
Lesson 2: Planning and Designing a Cloud Environment
Topic 2A: Meet Cloud Business Requirements
Topic 2B: Design Capacity Planning and Requirements
Lesson 3: Administering Cloud Resources
Topic 3A: Manage Cloud Administration46
Topic 3B: Manage Compute Resources in the Cloud
Topic 3C: Manage Memory Resources64
Lesson 4: Managing Cloud Storage
Topic 4A: Understand Cloud Storage Types68
Topic 4B: Configure Cloud Storage Solutions74
Topic 4C: Configure Cloud Storage Protocols and RAID
Lesson 5: Managing Networks in the Cloud91
Topic 5A: Deploying Cloud Network Services
Topic 5B: Identify Cloud Network Infrastructure Components
Lesson 6: Securing and Troubleshooting Networks in the Cloud
Topic 6A: Secure a Network in a Cloud Environment
Topic 6B: Troubleshooting Cloud Connectivity128
Lesson 7: Managing Cloud Migrations and Troubleshooting Cloud Deployments143
Topic 7A: Manage Cloud Migrations
Topic 7B: Troubleshoot Cloud Deployment and Migration Issues

Lesson 8: Managing Cloud Automation and Orchestration	163
Topic 8A: Understand Cloud Automation and Orchestration Techniques 1	164
Topic 8B: Troubleshoot Automation and Orchestration in the Cloud	175

Lesson 9: Understanding Cloud Security Concepts18	81
Topic 9A: Administer Identity and Access Management in the Cloud18	82
Topic 9B: Manage Cloud Operating System and Application Security 19	92
Topic 9C: Manage Data Security and Compliance in the Cloud	00

Lesson 10: Managing Cloud Security	. 213
Topic 10A: Implement Security Measures in the Cloud Domain	. 214
Topic 10B: Troubleshoot Cloud Security	. 224

Lesson 11: Managing Cloud Performance	237
Topic 11A: Operate Efficiently in the Cloud	238
Topic 11B: Accomplish Cloud Operations Tasks	242
Topic 11C: Optimize Cloud	249
Topic 11D: Troubleshoot Common Cloud Performance Problems	259

Lesson 12: Managing Maintenance in the Cloud267
Topic 12A: Configure Logs, Monitoring, and Alerting for Cloud Services 268
Topic 12B: Manage Backup and Restore in the Cloud

Lesson 13: Implementing High Availability and Disaster Recovery in the Cloud	287
Topic 13A: Understand High Availability and Scaling in the Cloud	288
Topic 13B: Manage Disaster Recovery in the Cloud	297
Topic 13C: Manage Backup and Restore in the Cloud	305

Appendix A: Mapping Course Content to CompTIA Cloud+ (CV0-003)	A-1
Solutions	S-1
Glossary	G-1
Index	. I-1

About This Course

CompTIA is a not-for-profit trade association with the purpose of advancing the interests of IT professionals and IT channel organizations; its industry-leading IT certifications are an important part of that mission. CompTIA's Cloud+ Certification is an upper-level certification designed for professionals with 2–3 years of hands-on work experience in a systems administrator job role.

This exam will certify the successful candidate has the knowledge and skills required to understand cloud architecture and design, deploy cloud services and solutions, successfully maintain, secure, and optimize a cloud environment, and to troubleshoot common cloud management issues.

CompTIA Cloud+ Exam Objectives Blueprint

Course Description

Course Objectives

This course can benefit you in two ways. If you intend to pass the CompTIA Cloud+ (Exam CV0-003) certification examination, this course can be a significant part of your preparation. But certification is not the only key to professional success in the field of cloud management. Today's job market demands individuals with demonstrable skills, and the information and activities in this course can help you build your cloud sysadmin skill set so that you can confidently perform your duties in any cloud administrator role.

On course completion, you will be able to:

- Understand cloud concepts
- Plan and design a cloud environment
- Administer cloud resources
- Manage cloud storage
- Manage networks in the cloud
- Secure and troubleshoot networks in the cloud
- · Manage cloud migrations and troubleshoot cloud deployments
- Manage cloud automation and orchestration
- Understand cloud security concepts
- Manage cloud security
- Manage cloud performance
- Manage cloud maintenance
- Implement high availability and disaster recovery in the cloud

Target Student

The Official CompTIA Cloud+ *Guide (Exam CV0-003)* is the primary course you will need to take if your job responsibilities include cloud architecture and deployment, optimization, and cloud security within your organization. You can take this course to prepare for the CompTIA Cloud+ (Exam CV0-003) certification examination.

Prerequisites

To ensure your success in this course, you should have systems administration skills comprising 2–3 years' experience. CompTIA Network+ and Server+ certification, or the equivalent knowledge, is strongly recommended.



The prerequisites for this course might differ significantly from the prerequisites for the CompTIA certification exams. For the most up-to-date information about the exam prerequisites, complete the form on this page: www.comptia.org/training/resources/exam-objectives

How to Use the Study Notes

The following notes will help you understand how the course structure and components are designed to support mastery of the competencies and tasks associated with the target job roles and help you to prepare to take the certification exam.

As You Learn

At the top level, this course is divided into **lessons**, each representing an area of competency within the target job roles. Each lesson is composed of a number of topics. A **topic** contains subjects that are related to a discrete job task, mapped to objectives and content examples in the CompTIA exam objectives document. Rather than follow the exam domains and objectives sequence, lessons and topics are arranged in order of increasing proficiency. Each topic is intended to be studied within a short period (typically 30 minutes at most). Each topic is concluded by one or more activities, designed to help you apply your understanding of the study notes to practical scenarios and tasks.

In addition to the study content in the lessons, there is a glossary of the terms and concepts used throughout the course. There is also an index to assist in locating particular terminology, concepts, technologies, and tasks within the lesson and topic content.



In many electronic versions of the book, you can click links on key words in the topic content to move to the associated glossary definition, and on page references in the index to move to that term in the content. To return to the previous location in the document after clicking a link, use the appropriate functionality in your eBook viewing software.

Watch throughout the material for the following visual cues.

Student Icon	Student Icon Descriptive Text
	A Note provides additional information, guidance, or hints about a topic or task.
A	A Caution note makes you aware of places where you need to be particularly careful with your actions, settings, or decisions so that you can be sure to get the desired results of an activity or task.

As You Review

Any method of instruction is only as effective as the time and effort you, the student, are willing to invest in it. In addition, some of the information that you learn in class may not be important to you immediately, but it may become important later. For this reason, we encourage you to spend some time reviewing the content of the course after your time in the classroom.

Following the lesson content, you will find a table mapping the lessons and topics to the exam domains, objectives, and content examples. You can use this as a checklist as you prepare to take the exam, and review any content that you are uncertain about.

As a Reference

The organization and layout of this book make it an easy-to-use resource for future reference. Guidelines can be used during class and as after-class references when you're back on the job and need to refresh your understanding. Taking advantage of the glossary, index, and table of contents, you can use this book as a first source of definitions, background information, and summaries.

Lesson 1

Understanding Cloud Concepts

LESSON INTRODUCTION

In order to better understand cloud services and providers, a consistent set of concepts must be applied. This lesson covers cloud concepts and terminology. Cloud service providers have their own unique names for their offerings, so it is easy to get confused. The final topic in the lesson addresses a common troubleshooting methodology that applies to on-premises and cloud deployments.

Lesson Objectives

In this lesson, you will:

- Understand cloud concepts.
- Understand cloud terms and concepts.
- Understand the troubleshooting methodology.

Topic 1A

Recognize Cloud Concepts



EXAM OBJECTIVES COVERED

1.1 Compare and contrast the different types of cloud models

The term "cloud" can have various meanings. In order to reduce confusion, the NIST definition of cloud services will be used. The three primary cloud service models are discussed next, followed by the cloud deployment models. Cloud services and the shared security model conclude this section.

What Is the Cloud?

The definition of "cloud computing" may be confusing. There are many different offerings, with different names, that are provided by different vendors. Five characteristics have been agreed upon for a basic understanding of what cloud computing means. These characteristics help to illustrate cloud services. The five characteristics defined by the **National Institute of Standards and Technology** (NIST) are:

- **On-demand self-service:** Consumers can provision resources as needed and automatically.
- **Broad network access:** Services are available across the network from commonly available clients.
- **Resource pooling:** The cloud service provider (CSP) pools resources in a multitenant model and adjusts resource allocation on an on-demand basis, and the specific distribution of hardware resources is abstracted from the consumer.
- **Rapid elasticity:** Resources are provisioned and released to adjust for changes in demand and consumption. This process may be automatic or manual.
- **Measured service:** Metering of resources is monitored, controlled, and billable.



You will find the remainder of this class much easier if you memorize these five characteristics quickly.

On-Demand Self-Service

Cloud services consumers can provision services on an as-needed basis, without the need to work with the **CSP** directly. These resources might include additional compute power, additional storage, new websites, or even database services. The consumer can expand (or reduce) these services without the need for human assistance from the CSP.

Broad Network Access

Client devices and traditional server deployments are able to access cloud-based resources across the network. The network might include the local on-premises network or the Internet, or both. Cloud resources have the potential to be globally accessible.

Resource Pooling

CSPs pool network, storage, and compute capabilities and then dynamically and automatically allocate those resources to consumers on an on-demand basis. The consumers do not know (or care) where the resources might physically be located. The next time those services are used by the consumer, the resource locations might have changed. The CSP manages the resources and maximizes their use.

Rapid Elasticity

Server resources in a traditional model are purchased as a capital expenditure, and whether or not those resources are efficiently utilized, their cost and capabilities are fixed. In some business models, resource needs change throughout the year. For example, retail demand is significantly higher during some parts of the year than others. With cloud-based computing, resources are dynamically allocated, making for far more efficient utilization of those resources. Servers that might have been under utilized for most of the year no longer need to be purchased and maintained.

Measured Service

CSPs meter the utilization of their resources. This metering permits more efficient and dynamic resource allocation. It also permits the CSPs to bill consumers accurately for exactly the quantity of resources consumed.

Understand Cloud Services Models

One of the simplest explanations of cloud services is the idea of offloading responsibility. In a traditional client-server model, companies do not offload responsibility for their servers, storage, or networking. Instead, they purchase expensive, complex, and often proprietary hardware, hire a team of experts to manage that hardware, and cannot easily scale their investment as their business needs change. With the cloud, some or all of those responsibilities can be offloaded to a CSP.

Cloud service models are initially divided into three types of solutions. These may be differentiated by what responsibilities are being offloaded. Many more cloud service models have since been defined, but these three best exemplify the primary aspects of cloud computing:

- **Software as a Service (SaaS):** The consumer is being provided with the direct use of the software. Responsibility for the hardware where that software runs, the operating system upon which it runs, and the installation and patching of the software itself are all offloaded to the CSP.
- **Platform as a Service (PaaS):** The service structure is provided by the CSP. It is up to the consumer to populate that structure, manage it on a day-to-day basis, and assume responsibility for the content. Support for the hardware, as well as the service platform that hosts the content, is offloaded to the CSP.
- **Infrastructure as a Service (IaaS):** The hardware infrastructure is provided to the consumer. The consumer assumes responsibility for all layers above that hardware. The CSP manages hardware failures, firmware updates, device drivers, and hardware compatibility. The consumer installs and manages the operating system on top of the hardware as well as any services and applications that run above that operating system.

• **Anything as a Service (XaaS):** "Anything as a service" is a catch-all phrase for technology solutions that are moved to the cloud. There are many IT-oriented examples, including database as a service (DBaaS), desktop as a service (DaaS), and containers as a service (CaaS).



Areas of responsibility are depicted for each cloud service model as compared to a traditional IT deployment. (Image © 123rf.com.)

Read carefully! This section contains critical information and terms that are used extensively in the course.

Software as a Service

SaaS permits consumers to use the software provided by the CSP. The CSP retains responsibility for installing, configuring, maintaining, patching, and upgrading the software. The software is typically accessible from many client device platforms, such as phones, tablets, and traditional computers.

The software is licensed for use by the consumer. Licensing is usually based on a subscription model, where only the number of deployments needed is purchased and paid for. SaaS licensing helps to make the cost management of software more streamlined and scalable. SaaS enjoys a very high adoption rate on the Internet. The consumer may realize lower initial deployment costs, quicker deployments, and lower total cost of ownership expenditures over the software lifecycle by offloading the support and maintenance of the software to a service provider.

SaaS is also operating system agnostic, meaning that organizations will rarely have to worry about the preferred operating system platform of a given employee. The cloud-based software will also likely be available no matter what hardware platform employees might use. For example, in your organization you might have fifty employees using Microsoft Windows, fifty using Apple macOS, and fifty more using Ubuntu Linux. Your organization could deploy a SaaS solution such as Dropbox (cloud-based storage) that supports all three operating systems. In a traditional non-cloud implementation, it would be much harder to provide a centralized, policy-managed, and scalable storage solution that supports all three operating systems.

SaaS examples:

- Microsoft Office 365
- Google Apps
- WebEx
- Dropbox
- Netflix

Target audience:

• End users



SaaS provides hosted applications and all of the corresponding software and hardware infrastructure related to using that application. (Image © 123RF.com.)

Platform as a Service

PaaS scenarios provide the hardware, operating system, and necessary tools to consumers. The consumers then utilize the tools to manage their data on their own. The CSP is responsible for hardware support and operating system support as well as platform maintenance. The consumer simply uses the platform within the scope of their own business needs.

PaaS solutions are often aimed at developers and database administrators (DBAs). These individuals use the provided platform to develop whatever applications or database services are needed by the organization without having to first build the platforms. PaaS solutions also scale quickly and easily, providing consistent development platforms as needs change.

By offloading responsibility for the hardware and operating system to the CSP, developers do not have to concern themselves with supporting the platform they are working on. Menial tasks, such as OS and application updates, hardware failures, and device drivers, are no longer a concern. For example, a DBA no longer has to request a server, install an operating system, and then install a database platform such as Microsoft SQL Server or MariaDB. Those tasks have all been offloaded to the CSP.

PaaS solutions typically support multiple development environments and programming languages. They also support development for all platforms, including phones, tablets, end-user workstations, and servers.

PaaS examples:

- Google App Engine
- Heroku
- AWS ElasticBeanstalk
- Salesforce

Target audience:

- Developers
- DBAs



PaaS provides infrastructure software, operating systems, virtualization, servers and server hardware, networking, and data center electrical and mechanical operations. (Image © 123RF.com.)

Infrastructure as a Service

laaS offloads responsibility for hardware support to the CSP. Consumers are responsible for the management of virtual machines hosted on the CSP hardware infrastructure. The consumers install operating systems themselves, such as Microsoft Windows or Red Hat Enterprise Linux (RHEL), configure and patch the OSs, and install software. laaS can be much more cost effective for consumers because they can deploy exactly the server platforms they need. laaS also provides easy scalability. This may result in a significant reduction in capital expenditures for hardware and licenses. In a traditional infrastructure model, companies are required to predict server needs during a budget cycle and estimate the overall utilization of the services. Such a model makes it very difficult to be agile or efficiently utilize services in industries in which demand varies significantly throughout the year (such as retail holiday shopping). The pay-as-you-go model of laaS means that organizations can pay for what they need—no more and no less.

Onsite servers also require a great deal of infrastructure support. Servers consume a lot of physical space, must be properly and reliably powered, must be continually cooled, and must be physically secure. All of these requirements mean significant operating expenditures for businesses. By offloading these costs to a CSP, businesses may save a great deal of money and time.

laaS examples:

- AWS EC2
- Microsoft Azure
- Rackspace
- Digital Ocean

Target audience:

• IT administrators



laaS provides virtualization, servers and server hardware, networking, and data center electrical and mechanical operations to clients. (Image © 123RF.com.)

Anything as a Service

XaaS, where "X" represents any possible service, began as a way of taking the agility, speed, and reduced initial capital expenditures that occur with the other cloudbased "as a service" deployments and shifting those benefits to other IT services. These services might include email, desktop operating systems, remote access, and even security. Many innovative companies have creatively carved a niche within the "as a service" role for themselves. XaaS means any delivery by the Internet with a flexible, pay-as-you-go structure. This new service model provides reduced capital expenditures and business agility.

Understand Cloud Deployment Models

Cloud Components and Clients

There are three main components in a cloud services solution. The first component is the client platform from which the cloud services are being accessed. The second is the data center where the cloud services are being hosted. The final component is the network connection between those two points.



The three components of cloud services are the client device, the data center, and the network that links them together. (Images © 123RF.com)

Cloud Service Component	Role
CSP data center	Hosts cloud services
Client	Means of access to cloud services for consumer
Network	Path between cloud services and client devices
Clo	ud services components.

Major CSPs, such as Microsoft and Amazon, have a great many data centers distributed across the world. These data centers are redundant, have extremely reliable access to power, have extremely reliable Internet access, and are physically secure. Cloud services are hosted within the walls of these data centers.

Cloud services consumers could be virtually anyone, on any platform. Cloud services may include storage, email, e-commerce, office suites, and development environments. Users may access these services from phones, tablets, traditional computers, Internet of Things (IoT) devices, and servers. The cloud client devices may be any device with a network connection. The major operating systems on the client devices include Microsoft Windows, Apple macOS, Linux, iOS, and Android.

The network is the path between the CSP data centers and the client devices. In some deployment models, the network connection may be wholly owned and operated by your company. In other cases, the Internet may be the network path to cloud services. Access may also come via cell connections. In some cases, all three network connection types may be used.

Understand Cloud Deployment Models

The cloud service infrastructure can be managed internally by a single organization for its own use or managed by a CSP that provides services to many organizations. These two models can be combined into a hybrid solution. A community of consumers can also choose to implement a cloud hosting solution.

- **Public cloud:** A CSP owns the cloud deployment and allocates its resources to external, unaffiliated customers. Those customers share the public cloud's resources without knowing precisely where their data is in relation to that of any other organization.
- Private cloud: Services are provided to only a single organization.
- **Community cloud:** Services are offered to several organizations that may have similar service needs but are otherwise autonomous.
- **Hybrid cloud:** There is a combination of two or more private, public, or community deployments.

Public Cloud



Public cloud environments are shared among many unrelated organizations. (Images © 123RF.com.)

CSPs offer public cloud services to virtually any customer. Customers use a subscription model to pay for access. The resources provided by the CSP are then shared among its customers. No customer has any real understanding of precisely where their resources or data may be at a given moment. In the background, the CSP dynamically reallocates resources throughout the data center to support the current demand.

Public cloud services are what most people think of when they hear about "cloud computing." Amazon and Microsoft are two major public cloud vendors (though they each also offer private cloud deployment options).

Private Cloud



A private cloud model is available only to the enterprise that owns it. (Image © 123RF.com.)

Some organizations may choose to implement their own internal cloud solution. The company itself will provide a data center and virtualization and offer a catalog of available cloud services to the rest of the organization. The company retains complete control of the cloud deployment but can leverage the advantages of cloud technologies. This is especially viable for very large enterprises.

In some cases, a private cloud may be necessary for security concerns. An organization may be unable to use the public cloud deployment model due to industry regulations, government requirements, insurance, or other reasons.



Community Cloud

of related organizations. (Images © 123RF.com.)

A community cloud exists to serve a discrete body of consumers who have similar business or security needs. Access to this cloud deployment model is limited to the members of the community. The community cloud may be owned or managed by any one or more of the community members, or it may be managed by a third party.

Hybrid Cloud



Hybrid clouds are any combination of the other three models. (Image © 123RF.com.)

A hybrid cloud deployment is any combination of the other three deployment models. For example, an organization may choose to utilize some services offered via a CSP's public cloud while hosting other services in a private cloud environment. The services in the public cloud portion may be cheaper, and security may be less of a concern. The services hosted in the private cloud may be more secure, but the deployment is more expensive.

Cloud Within a Cloud

Also known as a **virtual private cloud (VPC)**, the concept of cloud within a cloud means a public CSP hosts your organization's cloud services in an isolated segment, separated from any resources shared with other companies. While your cloud services exist within a public cloud, those resources are private and unshared. Your organization takes full administrative responsibility.

Virtual private clouds are an example of single-tenant deployments.

Sometimes there is confusion regarding the difference between VPC and private cloud deployments. Here are some key differences:

- VPC–logical isolation of the cloud deployment that resides on a CSP's infrastructure (and is therefore scalable).
- Private cloud–physical isolation of the cloud deployment in a private data center, a community data center, or CSP's infrastructure. It is limited by the available hardware and therefore less scalable.

Multitenancy

Multitenancy is the concept behind public cloud deployments. Multiple consumers, known as tenants, share computing resources owned and managed by the CSP. This is the opposite idea from a VPC deployment. It is multitenancy that provides the cost benefits behind shared resource utilization.

Multi-cloud

There are many **multi-cloud** variations, but some of the most common are combinations of cloud services spread among two or more public CSPs (such as AWS and Azure) as well as on a private cloud infrastructure.

Multi-cloud deployments reduce reliance on a single vendor, provide greater service flexibility and choice, permit improved geographic control of data, and help manage disaster mitigation.



private cloud platforms. (Images © 123RF.com)

Recognize Cloud Services

Google Workspace

Formerly known as G Suite and Google Apps, Google Workspace provides word processing, spreadsheets, presentation software, and other services. Google Workspace is an example of SaaS.

Office 365

Microsoft's Office suite is available online as a SaaS product. Office 365 provides a great deal of flexibility, platform independence, and ease of installation and support.

Digital Ocean

Digital Ocean is a developer-oriented cloud service provider that offers scalable and quickly deployable resources. This very popular web application hosting service is known for simplicity.

Rackspace

Rackspace provides cloud servers, database platforms, load balancers, storage, and other services to organizations. One thing that Rackspace is well known for is "fanatical" support.

Red Hat Cloud Suite

Red Hat, originally known for its enterprise Linux operating system and supporting services, offers Red Hat Cloud Suite for cloud services. The suite consists of four key products: OpenStack Platform (for building public and private clouds), Virtualization (for virtualizing cloud-based servers), Satellite (for cloud services management), and OpenShift (for Kubernetes container management).

Recognize Advanced Cloud Services

Internet of Things

The **Internet of Things (IoT)** refers to a combination of network connectivity and smart devices that facilitate the collection and analysis of data. These devices may include software, sensors, and robotics that exchange data and instructions over the Internet or internal networks. The IoT is enabled by nearly global network connectivity, low-cost sensors to collect data, and cloud management platforms.

Common uses for IoT products include:

- Smart homes
- Medical monitoring
- Agriculture management
- Energy management
- Manufacturing/Industrial production

Serverless Computing

Serverless computing still utilizes compute resources, contrary to what the name implies. Compute resources are allocated on demand to applications, and no resources are reserved when the application is not in use. Billing reflects the application's actual use of resources. Serverless environments require no configuration, monitoring, or capacity planning.

Serverless computing is also known as Function as a Service (FaaS).

Machine Learning and Artificial Intelligence (AI)

Artificial intelligence (AI) is concerned with simulating human intelligence by providing structured, semi-structured, and unstructured data and solving complex problems. Al accomplishes this by using a set of rules to manage its analysis.

Machine Learning (ML) is a subset of AI. The goal of ML is to make accurate predictions by extracting data based on learned information and experience. ML systems are not explicitly programmed to find a particular outcome. Instead, they are programmed to learn from provided data and then make accurate decisions based on what they've learned. Insights are gained with minimal human interaction.

Deep Learning (DL) is a subset of ML. DL provides a greater degree of accuracy when analyzing unstructured data.

Understand the Largest Cloud Service Providers

The following three CSPs make up the bulk of the global cloud service offerings. All three are in direct competition with each other and therefore offer many similar services, but with different names. All three CSPs provide a vast number of cloud offerings via global infrastructures.

Amazon Web Services

Amazon Web Services (AWS) offers around 200 cloud-based services supported by a global infrastructure of data centers that consists of 77+ availability zones and 24+ regions. In keeping with the cloud services attributes, AWS allows for quick provisioning and low up-front costs.

AWS is built primarily on a Linux infrastructure.

Some of the top AWS services include:

- Simple Storage Service (S3): storage as a service
- Elastic Compute Cloud (EC2): infrastructure as a service
- Lambda: serverless compute services
- Glacier: long-term storage and archiving service
- Simple Notification Service (SNS): service for publishers to push messages to subscribers
- CloudFront: dynamic web site services

Microsoft Azure

Microsoft Azure also presents approximately 200 cloud products. It too has a globally dispersed infrastructure of data centers and networks divided into various regions and availability zones. Azure permits integration of services, quick deployments, and low capital expenditures to get started. It is built on a mixture of Windows and Linux solutions.

Some of the top Azure services include:

- Azure Virtual Machines: IaaS
- Azure Disk Storage: STaaS
- Azure Visual Studio: development environment
- Azure Functions: serverless compute service
- Azure Backup: backup and restore service
- Azure SQL: SQL database service
- Azure Cosmos DB: NoSQL database service
- Azure Active Directory: identity and access management service

Google Cloud Platform

Google Cloud Platform (GCP) offers a large array of cloud services spread across a global infrastructure consisting of data centers and network connectivity. The infrastructure includes 73+ zones and 24+ regions and continues to grow. Like AWS and Azure, GCP leverages the quick deployments and low up-front costs of cloud services to offer compute, storage, and database services.

Some of the top GCP services include:

- Cloud Storage: STaaS
- Compute Engine: laaS
- App Engine: PaaS

- Cloud SQL: SQL database service
- Firestore: NoSQL database service
- Big Query: big data services

Understand the Role of Cloud Managed Service Providers

It can be difficult for organizations—particularly small and medium-sized companies to provide the broad range of expertise needed for every aspect of cloud computing.

There are also third-party **managed service providers (MSPs)** that are independent of the CSP. Your organization may choose to outsource cloud design, migration, deployment, and management solutions to these companies, relying on their expertise and experience. Many CSPs also offer management services for their products. For example, AWS offers AWS Managed Services.

Such outsourcing may result in reduced operating expenditures and greater efficiency.

Some services provided include:

- Reporting and monitoring
- Backup and recovery processes
- Performance testing

Some benefits of MSPs include:

- Expertise beyond what your in-house staff can achieve (such as cybersecurity, disaster recovery, compliance, and future technologies).
- Freed time for your in-house IT and development staff to spend on other projects.



Using an MSP outsources many business tasks. (Images © 123RF.com.)



You may wish to discuss whether or not your business has considered working with an MSP for cloud services.

Understand the Shared Security Model

In the shared responsibility security model, CSPs and consumers manage different aspects of cloud security. CSPs will typically provide physical security for the data center as well as isolate data between customers. Cloud consumers are responsible for the direct user access security for their own data. In other words, the CSP secures the data center where files are stored, but the consumer secures the files themselves by using permissions and encryption. AWS phrases it this way: AWS is responsible for the security *of* the cloud; the consumer is responsible for security *in* the cloud.



The shared security model divides responsibility between the CSP and the consumer.



The AWS shared responsibility model: <u>https://aws.amazon.com/compliance/shared-</u> <i>responsibility-model/

Review Activity:

Cloud Concepts

Answer the following questions:

- 1. You are helping to design a cloud architecture for your organization. Industry regulations require the organization to keep all customer data on its own systems, rather than housing this data on a system shared with other organizations. Which cloud deployment model satisfies this requirement?
- 2. You are advising a small company on cloud deployment and management options. The company does not have a dedicated IT staff and does not intend to increase staff headcount. Recommend a solution.
- 3. Compare the areas of responsibility for the client organization and the hosting public cloud service provider (CSP) as related to the shared security model.

Topic 1B

Recognize Cloud Terms



EXAM OBJECTIVES COVERED

1.1 Compare and contrast the different types of cloud models

Implementing cloud services includes some specific tasks, no matter which service provider your organization selects. Concepts such as subscriptions, provisioning, virtual machines, containers, and scaling need to be understood and given context. This section covers those and other terms, providing a useful vocabulary for understanding cloud services.

Subscription Services

This payment model uses a recurring, periodic billing cycle that is often based on the length of the subscription (long-term subscriptions are often less expensive than short-term subscriptions). The model usually includes no long-term contracts. Access to services is provided as soon as the subscription is established and often can be terminated at any time. As an example, Microsoft Office 365 is purchased via a monthly or annual **subscription service**.

Identity Management

Identity management is the process by which identities are established and access to resources is controlled. Typically, users are identified and assigned a user account. The account is then assigned rights and restrictions, which are enforced by access control systems such as permissions.

User accounts may be organized into groups to make management easier, but the overall process related to a given user has a specified level of access to resources.



Identity management is also known as Identity and Access Management or IAM.

Provisioning

Provisioning is one of several steps in the cloud services deployment process. The term refers to the allocation of cloud resources in the overall enterprise infrastructure. The provisioning process is governed by objectives, policies, and procedures for deploying services and data.

Provisioning may be accomplished via web-based or command-line interfaces. Provisioning is usually self-service, reflecting one of the NIST cloud characteristics discussed earlier.

In the overall deployment process, provisioning occurs before server, service, user, or network configuration. Access controls are usually part of the provisioning process.

Applications

With cloud applications, the installation and processing occur in the cloud, rather than on local workstations or servers. The cloud may be a private or public network. The applications are accessed over the network. One advantage of cloud applications is a consistent experience for all users, whether they use the same workstation platform (Windows, Linux, macOS) or mobile device (Android or IOS).

Virtual Machines

Virtualization allocates hardware resources among one or more **virtual machines** (**VMs**). The VMs then have an operating system and one or more applications installed on them. The VM participates on the network as a regular node, providing database, authentication, storage, or other services. VMs have greater access to hardware resources and can be provided with redundancy to increase high availability.

VMs are a key component of cloud-based laaS services, such as AWS EC2 or Azure Virtual Machines.

Containers

Containerization is a form of virtualization, but it is significantly different than VMs. Containers virtualize at the OS layer, rather than the hardware layer. A container holds a single application and everything it needs to run. This narrow focus allows containers to excel with technologies such as microservices. Containers are very lightweight, share a single OS (usually Linux), and provide a single function. GCP, Azure, and AWS all offer cloud-based container services.



Containers for each application, sharing bins, libraries, and a single OS.

Templates

Virtual machines may be deployed using templates. **VM Templates** reduce the confusion, misconfiguration, and cost associated with manually created VMs by providing standardized VM configurations.

For example, in your organization's private cloud, you might permit developers to create their own VMs (the self-service cloud characteristic). Rather than a developer having to fumble through the configurations (potentially making an expensive mistake), you can provide a template that provides the right amount of compute and storage resources, as well as the appropriate operating system, applications, and security configurations. The developer selects the template, and a VM is created from the established settings.

CSPs also use templates to offer flexible but standardized VM configurations to customers.

Post-Deployment Validation

Post-deployment validation ensures that deployed apps or services meet required service levels. Depending on the service, this may be handled through regression or functionality testing. If possible, automate post-deployment validation for efficiency and consistency.

Auto-scaling

Auto-scaling takes advantage of automated deployments and virtualization to provide the appropriate resources for the current demand. Resources can be scaled up or down to manage costs. Your organization only pays for the resources that it consumes. Auto-scaling is useful when resource utilization is difficult to predict or is seasonal.

Resources may be scaled up (more compute power, such as RAM, given to a single virtual server) or scaled out (more virtual servers deployed). When demand is reduced, the resources are reduced, saving money.

Hyperconverged

Hyperconverged cloud solutions combine compute, storage, and network resources into a single component. The goal is to reduce complexity and increase scalability. With **hyperconvergence**, the compute, storage, and network resources are inseparable and managed as a single unit. This is in contrast to converged resources which can be broken out into their constituent parts (compute, storage, networking), making them more difficult to manage—sysadmins are managing three pieces rather than one. **Review Activity:**

Cloud Terms

Answer the following questions:

- 1. Your organization's workflow is moving toward increased agility and the microservices that support this agility. Which virtualization technology best supports this?
- 2. You are helping to design a cloud architecture for your organization. The CIO wants to ensure that resources can scale up to meeting demand with minimal human intervention. Explain how virtualization and automated deployments help you meet the CIO's requirements.
- 3. Contrast converged and hyperconverged cloud resources.

Topic 1C

Understand the Troubleshooting Methodology



EXAM OBJECTIVES COVERED

5.1 Given a scenario, use the troubleshooting methodology to resolve cloud issues

One of the primary skills and duties of a cloud administrator is to troubleshoot problems with cloud services and data access. It is important to have a methodology for troubleshooting. You should also recognize that troubleshooting methods may change by situation, by skill level, and by experience with the cloud environment.

Troubleshooting Methodology

The following list represents the basic steps in a troubleshooting methodology:

- Identify the problem.
- Determine the scope of the problem.
- Establish a theory of probable cause, or question the obvious.
- Test the theory to determine the cause.
- Establish a plan of action.
- Implement the solution, or escalate.
- Verify full system functionality.
- Implement preventive measures.
- Perform a root cause analysis.
- Document findings, actions, and outcomes throughout the process.

Let's examine these steps in more detail.

Identify the Problem

The first troubleshooting phase is to identify the problem. This problem may be discovered for you by the end-users you support, exposed by log files, identified by monitoring software, or indicated by alerts on dashboard interfaces. There are many ways in which the problem may be identified.

There is a common concept in troubleshooting: If it worked yesterday, and it doesn't work today, what changed? Investigating recent changes to the network or service infrastructures is a good place to begin troubleshooting. In addition, consider non-infrastructure and non-IT changes, such as environmental, location, or facility changes. Such changes might include changes to electrical systems, HVAC, or physical security.

Determine the Scope of the Problem

Once a problem has been identified, you may choose to gather additional information to determine the scope of the problem. You may start this process by asking users for additional details or by examining log files. It is also very useful to attempt to replicate the problem by asking users to show you what they were doing when the problem was encountered or to try to recreate the situation where the problem first arose.

It is a good practice to back up data if there is any risk to the data during the troubleshooting phase. You must use your own judgment as to whether a data backup is necessary before you begin troubleshooting. In some cases, data and services may be replicated throughout the CSP's environment, making backups less essential.

One of the most important steps is to determine whether the problem exists for only one user, or whether it exists for multiple users. In addition, determining whether cloud service interruptions exist at multiple locations or just a single location will help determine the scope of the problem. The problem could be network based, in which case multiple devices or locations may be affected. It could be software based, such as a misconfiguration. This too may impact multiple users or locations.

For example, if users at one branch office cannot access a particular web application, but users at all of the other branch offices can access the web app, then the problem's scope appears to be that isolated branch office. This indicates that the problem is probably not with the web application itself but rather somewhere between the web application and that branch office.

Establish a Theory of Probable Cause, or Question the Obvious

The next phase in troubleshooting is to establish a probable cause for the problem. It is essential to keep this step as simple as possible. Newer administrators may be tempted to believe that because cloud services are complex, the problem must also be complex. Troubleshooting often begins with very simple steps, such as confirming that a given service is started. More complex problems may require you to examine log files, talk to users or other administrators, or check hardware.

When troubleshooting, it is useful to identify any common elements or similar problems that might span multiple cloud services or web apps. Such common elements might include a recent software change or a new service configuration.

Once you define these common elements and understand the problem's scope, you can establish a theory or probable cause.

Test the Theory to Determine the Cause

Next, test the theory by verifying that the likely cause is indeed the culprit. This phase may involve research or additional testing. Very simple problems may actually be solved during this step.

If your theory is confirmed, then you will move on to the next phase, which is to establish a plan of action. If your theory is not confirmed, then you must establish and test a new theory.

Establish a Plan of Action

The plan of action for addressing the problem must recognize that service interruptions and data loss should be avoided. If a cloud-based virtual server needs to be brought down, or if data has been lost due to a hard disk drive failure, the end-users must be notified. In addition, the plan should include the steps to be taken. These steps should be defined ahead of time rather than created during the implementation of the solution. It is useful to provide the impacted users with an expected duration of the outage.

Implement the Solution, or Escalate

This phase involves following the plan of action established earlier. It is important to follow the plan of action and not to deviate. It is possible that you may not have the knowledge to implement the plan of action and will need to escalate the problem to the cloud vendor's support team, the cloud MSP or other members of your own team.

When following a plan of action, be sure to make only one change at a time, then test the result. If you make multiple changes simultaneously it becomes difficult to identify exactly which change corrected the problem. If a given change did not solve the problem, reverse that change, and then try another option.

Verify Full System Functionality

Once the potential solution has been implemented, the next phase is to test for functionality. Your goal is to ensure that the service or web application has returned to the established service levels.

Implement Preventive Measures

It may be possible to preemptively reconfigure other services or network devices to avoid a repeat of the problem. It may also be possible to implement additional technologies (such as RAID) or additional practices (such as backups) to prevent future instances of failure. Cloud services often implement scalability and high availability to ensure services are accessible to consumers.

Perform a Root Cause Analysis

Once service has been restored to your users, it is time to evaluate why the problem occurred. Identifying the root cause may permit you to avoid the problem in the future by changing processes or by implementing different technologies.

Document Findings, Actions, and Outcomes Throughout the Process

Documentation should be maintained throughout the service's lifecycle, including during the troubleshooting process. Documenting the symptoms of the problem, the results of research into potential solutions, and the results of each step of the plan of action (whether the step was successful or not) will permit you to better understand your environment and therefore help prevent possible future problems. Note that documentation is not a separate step but rather a good practice to be used during each phase of the troubleshooting process.



Some service desk management software requires the use of tickets. This software may require that troubleshooting documentation be entered before the ticket can be closed.

Understand the Environment, Including Corporate Policies, Their Procedures, and Impacts

Carefully consider corporate policies that govern troubleshooting and downtime. Standard operating procedures govern how particular tasks are to be accomplished, and such procedures should be followed during troubleshooting. Cloud policies govern how the organization interacts with cloud resources. In terms of troubleshooting, these policies manage case escalation (for example, escalation to the CSP's technical support). Such cases may incur a cost. Administrators must have a complete view of the cloud environment to understand troubleshooting before escalating cases upward.

In addition, service–level agreements (SLAs) might enforce penalties on your organization for service outages. Don't forget that cloud service providers also have SLAs backing their service availability. If an outage or other loss of access occurs, and it falls within the responsibility of the CSP, your organization may be awarded reduced service fees or other compensation.



Cloud policies are covered later in the course.

Review Activity:

Troubleshooting Methodology

Answer the following question:

- 1. You are assigned a help desk ticket to find out why a user cannot access cloud resource. Using the troubleshooting methodology, how would you narrow the scope of the problem?
- 2. You have reestablished access to cloud resources for a user by correcting a DNS entry. Next, you attempt to ascertain why the DNS entry was incorrect. Which troubleshooting step are you using?
- 3. Contrast service level agreements (SLA) and standard operating procedures (SOP).

Lesson 1

Summary

Cloud concepts, while fundamental, are important for a consistent understanding of terminology. In this lesson, you have examined cloud characteristics, deployment models, and services. You also examined the shared security model. Finally, you learned a useful and straightforward troubleshooting methodology that applies to on-premises and cloud environments.

- 1. Which of the five cloud characteristics will be the most helpful for your organization?
- 2. Which cloud services (SaaS, PaaS, IaaS) is your organizational already using?
- 3. Will your organization use the public, private, community, or hybrid cloud deployment model?
- 4. Which cloud services are you currently using for your personal services or data?