IBM Cloud Technical Advocate
Study Guide

This study guide will help prepare you for the
IBM Cloud Technical Advocate Certification Examination.

Thorough study is essential to a successful outcome on the exam.

- Clear your schedule.
- Find a quiet place to study.
- Focus on the content.

- Open the associated on-line course for reference.
- Locate the study guide.
- Download the study guide.

- Print a copy of the study guide.
- Take notes.
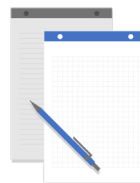
What's in the Study Guide

This study guide covers:

- ❑ **Introduction to Cloud**
- ❑ **IBM Cloud Fundamentals**
- ❑ **Networking**
- ❑ **Computing**

- ❑ **Storage**
- ❑ **Services**
- ❑ **Cloud Native Overview**

How to Use this Study Guide

Read the content.

Take notes.

Answer practice questions.

Courses and Objectives

Courses

1. Cloud Job Role Videos
2. Cloud Computing Characteristics, History, and Trends
3. Cloud Service and Deployment Models
4. Business Case Study for Cloud Computing

Objectives

- Recall the IBM Cloud job roles
- Define cloud computing and explain its essential characteristics, history, and emerging trends
- Identify service and deployment models of cloud computing and explain the benefits of each
- Identify key elements of business case studies for cloud computing

Course 1: Cloud Job Role Videos

Introduction and Objectives

In Course 1 of the study guide, the subject matter:

- Highlights a few of the different job roles that users of IBM Cloud have.

Lessons

- IBM Cloud Architect
- IBM Cloud Developer
- IBM Cloud Security Engineer
- IBM Cloud Site Reliability Engineer
- IBM Cloud Sales Engineer
- Summary

Objective

- Recall the IBM Cloud job roles

Course 1: Cloud Job Role Videos

IBM Cloud Job Roles

IBM Cloud Solution Architect	• Listen to client needs, understand and solve client issues, and collaborate with engineers and subject matter experts to design creative solutions for clients.
IBM Cloud Developer	• Take a conceptual design and bring it to life by building solutions and applications.
IBM Cloud Security Engineer	• Focus on detecting risk, protecting vulnerable resources, and keeping client data safe by building secure infrastructure, platforms, and software.
IBM Cloud Site Reliability Engineer	• Apply a specific skill set to analyze and learn from system issues to improve system availability and reliability.
IBM Cloud Sales Engineer	• Assist clients in creating action plans to help them resolve critical business problems. This job role requires both sales and technical skills.

Question 1.

Which IBM Cloud job role requires both sales and technical skills?

A.   IBM Cloud Architect

B.   IBM Cloud Site Reliability Engineer

C.   IBM Cloud Sales Engineer

D.   IBM Cloud Security Engineer

Answer C. The IBM Cloud Sales Engineer requires both sales and technical skills.

Question 2.

Which IBM Cloud job role focuses on detecting risk and keeping client and customer data safe?

A. IBM Cloud Developer

B. IBM Cloud Site Reliability Engineer

C. IBM Cloud Security Engineer

D. IBM Cloud Architect

Answer C. The IBM Cloud Security Engineer is focused on detecting risk and keeping client and customer data safe.

Question 3.

Which IBM Cloud job role can find root causes of outages, understands system issues, and is good at anticipating a client's potential points of failure in their solution?

A.   IBM Cloud Developer

B.   IBM Cloud Architect

C.   IBM Cloud Site Reliability Engineer

D.   IBM Cloud Security Engineer

Answer C. The IBM Cloud Site Reliability Engineer can find root causes of outages, understands system issues, and is good at anticipating a client's potential points of failure in their solution.

Question 4.

Which IBM Cloud job role listens to client needs, understands and solves client issues, and collaborates with engineers and subject matter experts to design solutions using IBM Cloud and other services?

A.   IBM Cloud Developer

B.   IBM Cloud Solution Architect

C.   IBM Cloud Site Reliability Engineer

D.   IBM Cloud Customer Success Manager

Answer B. The IBM Cloud Solution Architect listens to client needs, understands and solves client issues, and collaborates to design solutions using IBM Cloud and other services.

Course 2: Cloud Computing Characteristics, History, and Trends

Introduction and Objectives

In Course 2 of the study guide, the subject matter:

- Focuses on fundamental concepts of cloud computing.

- Highlights the high-level concepts of virtualization and hypervisors.

- Highlights the trends associated with cloud computing.

Lessons

- History of Cloud Computing
- Definition of Cloud Computing
- Characteristics of Cloud Computing
- High-Level Concepts of Virtualization and Hypervisors
- Trends associated with Cloud Computing
- Summary
- Knowledge Check

Objective

- Define cloud computing and explain its essential characteristics, history, and emerging trends

Course 2: Cloud Computing Characteristics, History, and Trends

The History of Cloud Computing

1950s

- The ability to have multiple users access and use a single machine dates back to the large mainframe systems that dominated the 1950s.

1960s

- The internet, as well as advances in computer's memory technology, began in the 1960s.
- In the beginning of the internet, IBM played a computing role in the Apollo missions, both on the ground and in the spacecraft.

1970s

- The concept of virtual machines (VMs) developed.
- Virtualization made it possible for multiple operating systems to function at the same time in the same environment.

1980s and 1990s

- Virtualized private networks offered the same service quality as dedicated services at lower costs.
- Hypervisors were enabled.

2000s to present

- Servers were virtualized into shared hosting environments, virtual private servers, and virtual dedicated servers, using the same types of functionality provided by the virtual machine operating system.
- Hypervisors enabled a company to take one physical node and split it into multiple virtual systems.
- The pay-as-you-go model was a key driver for cloud computing success.

Course 2: Cloud Computing Characteristics, History, and Trends

Definition of Cloud Computing

According to the US National Institute of Standards and Technology (NIST), cloud computing is defined as a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service provider interaction.

View the types of cloud computing and their definitions below.

Type of Cloud Computing	Definition
Public Cloud	A cloud service provider that makes computing resources available to users over the public internet for free, subscription based, or pay-per-usage pricing.
Private Cloud	All cloud infrastructure and computing resources are dedicated to and accessible by only one client.
Hybrid Cloud	This is a combination of public and private (on-premises) cloud environments. It connects private (on-premises) cloud services and public clouds into a single, flexible infrastructure.
Multicloud	This type uses two or more clouds from two or more different cloud providers.
Hybrid Multicloud	This type uses two or more public clouds together with a private cloud environment.

Course 2: Cloud Computing Characteristics, History, and Trends

Essential Cloud Computing Characteristics

- **On-Demand self service:** Computing resources are provisioned or de-provisioned as needed in an automated fashion without human intervention.

- **Broad network access:** Computing resources can be accessed from anywhere over a network using devices such as smartphones, tablets, laptops, and personal computers.

- **Resource pooling:** Computing resources are pooled to meet consumer demand so that resources can be assigned, reassigned, or redistributed.

- **Rapid elasticity:** Computing resources can be acquired or released based on demand. Cloud computing provides an illusion of infinite computing resources for the users.

- **Measured service:** Consumers only pay for services they use or reserve.

Definition of Virtualization

Virtualization enables the creation of multiple virtual computers on a single physical host computer. Each virtual guest has its own operating system and applications and behaves like a regular computer. The number of virtual machines that a physical computer can host depends on the physical resources available and the resource requirements of the virtual guests. Virtualization reduces IT costs because the virtual CPU resources are less than the actual physical CPU resources of the host computer.

Definition of Hypervisor

A hypervisor is a lightweight software layer. A hypervisor is an abstraction layer that abstracts physical compute resources, making them available virtually to virtual machines (VMs). The hypervisor assigns processors, memory, and storage to each VM. It keeps VMs separate, so that they do not interfere with each other.

Course 2: Cloud Computing Characteristics, History, and Trends

Advantages of Virtualization

✓ **Portability:** Virtual machines can be relocated among physical computers in a network. This means that they can be easily restarted on a different host in the event of a physical server failure.

✓ **Flexibility:** Creating a virtual machine is fast and easy because it can be cloned with the operating system already installed. They can be scaled up and down as required. They allow multiple virtual machines to run concurrently.

✓ **Security:** A virtual machine helps improve security because it is a file that can be scanned for malicious software by an external program. A snapshot of the virtual machine can be taken at any point in time and then restored to that state if it becomes infected with malware.

✓ **Resource Utilization and Improved ROI:** Virtualization reduces IT overheads by making better use of compute resources. A single server running an application may only use a fraction of the server's resources. Virtualization means that multiple servers can share the same hardware, making fuller use of its resources.

✓ **Scale:** Virtualization allows architects to scale applications more easily. Applications can be scaled vertically by adding more virtual CPU and memory to a particular virtual machine through the hypervisor. They can also be easily scaled horizontally by creating more virtual machines through cloning.

Cloud Computing Trends

Hybrid multicloud is emerging as a new trend for application deployment. This has been made possible through container technology, such as Red Hat® OpenShift®, which was an emerging technology just a few years ago.

Question 1.

The _____ organization defines cloud computing as a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service provider interaction.

A.   Federal Cloud Standards Board (FCSB)

B.   The Defense Advanced Research Projects Agency (DARPA)

C.   US National Institute of Standards and Technology (NIST)

D.   Cloud Native Computing Foundation (CNCF)

Answer C. US National Institute of Standards and Technology (NIST)

Question 2.

The practice of providing shared access to a single mainframe emerged in which decade?

A.   1940s

B.   1950s

C.   1960s

D.   1970s

Answer A. In the 1950s is when the practice of providing shared access to a single mainframe emerged.

Question 3.

All cloud infrastructure and computing resources are dedicated to and accessible by only one client is the definition of which type of cloud computing?

A. Hybrid cloud

B. Multicloud

C. Private cloud

D. Public cloud

Answer C. A private cloud makes computing resources available to users over the public internet.

Question 4.

On-demand self service and rapid elasticity are two essential characteristics of cloud computing. What is another essential characteristic?

A.   Limited network access

B.   Dedicated fixed resources

C.   Unmeasured services

D.   Resource pooling

Answer D. On-demand self service, broad network access, resource pooling, rapid elasticity, and measured services are all essential characteristics of cloud computing.

Question 5.

Allocates processors, memory, and _____ are functions of a hypervisor.

A.   Storage to cloud services

B.   Storage to databases

C.   Storage to containers

D.   Storage to virtual machines

Answer D. Allocates processors, memory, and storage to virtual machines are functions of a hypervisor.

Question 6.

Virtualization enables the creation of multiple virtual _____ on a single physical host computer and each virtual guest _____ operating system and applications.

A. Computers / has its own

B. Computers / shares

C. Hypervisors / has its own

D. Hypervisors / shares

Answer A. Virtualization enables the creation of multiple virtual computers on a single physical host computer and each virtual guest has its own operating system and applications.

Question 7.

Portability in virtualization is defined as: _____ can be relocated among physical computers in a network so that they can be easily restarted on a different host in the event of a physical server failure.

A.   Applications

B.   Databases

C.   Virtual machines

D.   Networks

Answer C. Virtual machines can be relocated among physical computers in a network so that they can be easily restarted on a different host in the event of a physical server failure.

Question 8.

Scalability as it relates to IBM Cloud workloads is described by which of the following?

A. Horizontally by adding more resources and bandwidth or vertically by creating more pods through cloning

B. Vertically by adding more virtual CPU and memory to a particular virtual machine or horizontally by creating more virtual machines through cloning

C. Vertically by adding more network bandwidth or horizontally by creating more virtual machines through cloning

D. Horizontally by adding more virtual CPU and memory to a particular virtual machine or vertically by creating more virtual machines through cloning

Answer B. Vertically by adding more virtual CPU and memory to a particular virtual machine or horizontally by creating more virtual machines through cloning

Question 9.

Container technology is an emerging trend in cloud computing focused on _____as a result of historical events or natural disasters.

A.   Private cloud

B.   Public cloud

C.   Virtual private cloud

D.   Hybrid multicloud

Answer D. Container technology is an emerging trend in cloud computing focused on hybrid multicloud as a result of historical events or natural disasters.

Course 3: Cloud Service and Deployment Models

Introduction and Objectives

In Course 3 of the study guide, the subject matter:

- Introduces the fundamental concepts and benefits of service and deployment models of cloud computing.

Lessons

- Cloud Service Models
- Cloud Deployment Models
- Summary
- Knowledge Check

Objective

- Identify the service and deployment models of cloud computing

Course 3: Cloud Service and Deployment Models

Common Cloud Service Models

The most common and popular types of cloud services offered are IaaS, PaaS, and SaaS. The cloud service provider owns, manages, and maintains the IT assets. Organizations access them with an internet connection and only pay for usage on a subscription basis.

IaaS

- Infrastructure as a Service
- Provides on-demand access to computing services, networks, bare metal servers, and storage.
- Can scale resources depending on need.

PaaS

- Platform as a Service
- Provides on-demand access to a complete cloud hosted platform, allowing developers to code applications using software services and tools.
- Cloud service provider handles infrastructure management, patches, and updates.

SaaS

- Software as a Service
- Offers a model for distributing software.
- Cloud providers host and maintain a running application for end users.

Course 3: Cloud Service and Deployment Models

Types of Cloud Deployment Models and Their Benefits

Public

- Owned by cloud provider and usage is shared by multiple consumers and organizations
- Offers the most significant savings
- Benefits: Unlimited access to certain resources through cloud providers and easy management of unexpected traffic surges

Private

- Exclusive use by a single organization
- Modernizes and unifies an organization's legacy and in-house systems
- Benefits: Greater ability to customize applications and the infrastructure, workloads are run behind an organization's firewall, offers greater control and security, compliance with industry or government regulations is simplified

Hybrid

- Uses on-premises private cloud and a third-party public cloud
- Benefits: Developer productivity increases, greater control over resources, security and compliance are implemented consistently, business operations are accelerated

Multicloud

- Use of cloud services from two or more vendors
- Benefits: Organizations can choose cloud services and technologies for different cloud providers and vendors, reduced vulnerability to outages, exposure to licensing, security, and compatibility

Note that the NIST defines four cloud deployment models: Public cloud, Private cloud, Hybrid cloud, and Community cloud.

Question 1.

The _____ deployment model could require all infrastructure to be purchased and installed by the client.

A.   Private cloud

B.   Multicloud

C.   Public cloud

D.   Hybrid cloud

Answer A. A private cloud deployment model could require all infrastructure to be purchased and installed by the client.

Question 2.

One advantage that a company can have by using the public cloud deployment model to deploy a new application is that there's no infrastructure to manage. What is another advantage?

A. Less scalable

B. Minimal investment

C. Skill neutral

D. Management complexity

Answer B. Minimal investment is another advantage a company can have by using the public cloud deployment model to deploy a new application.

Question 3.

An organization would like to take advantage of the security and control of a private cloud while also realizing the financial benefits of using the public cloud. Which cloud deployment model would fit their needs?

A.   Multicloud

B.   Mixed cloud

C.   Blended cloud

D.   Hybrid cloud

Answer D. Hybrid cloud allows an organization to take advantage of the security and control of a private cloud while also realizing the financial benefits of using the public cloud.

Question 4.

The _____ deployment model has the lowest up-front costs and complexity.

A.   Private cloud

B.   Hybrid cloud

C.   Multicloud

D.   Public cloud

Answer D. The public cloud deployment model has the lowest up-front costs.

Question 5.

Which cloud service model provides on-demand access to computing services, bare metal servers, networking, and storage?

A.   Platform as a Service

B.   Software as a Service

C.   Functions as a Service

D.   Infrastructure as a Service

Answer D. The Infrastructure as a Service provides on-demand access to computing services, bare metal servers, networking, and storage.

Question 6.

The cloud provider takes care of all infrastructure and application management in which of the following cloud service models?

- A.   Platform as a Service
- B.   Software as a Service
- C.   Application as a Service
- D.   Infrastructure as a Service

Answer B. Software as a Service takes care of all infrastructure and application management.

Question 7.

A new company would like to code their applications using software and tools hosted on IBM Cloud. Which cloud service model meets their requirements?

A.   Platform as a Service

B.   Software as a Service

C.   Toolchain as a Service

D.   Infrastructure as a Service

Answer A. Platform as a Service would allow a new company to code their applications using software and tools hosted on IBM Cloud.

Question 8.

What is a benefit for a client who is purchasing Software as a Service on IBM Cloud?

A. They can build, test, deploy, run, and update applications quickly.

B. They can order redundant servers.

C. They can scale computing resources up or down as needed.

D. All infrastructure and application management is handled by Software as a Service vendor.

Answer D. The benefit to a client who purchases Software as a Service is all infrastructure and application management is handled by Software as a Service vendor.

Course 4: Business Case Study for Cloud Computing

Introduction and Objectives

In Course 4 of the study guide, the subject matter:

- Provides an example of how a business might transform and modernize as a result of cloud computing changes.

Lessons

- BNP Paribas Case Study – Transformation and Modernization
- Summary
- Knowledge Check

Objective

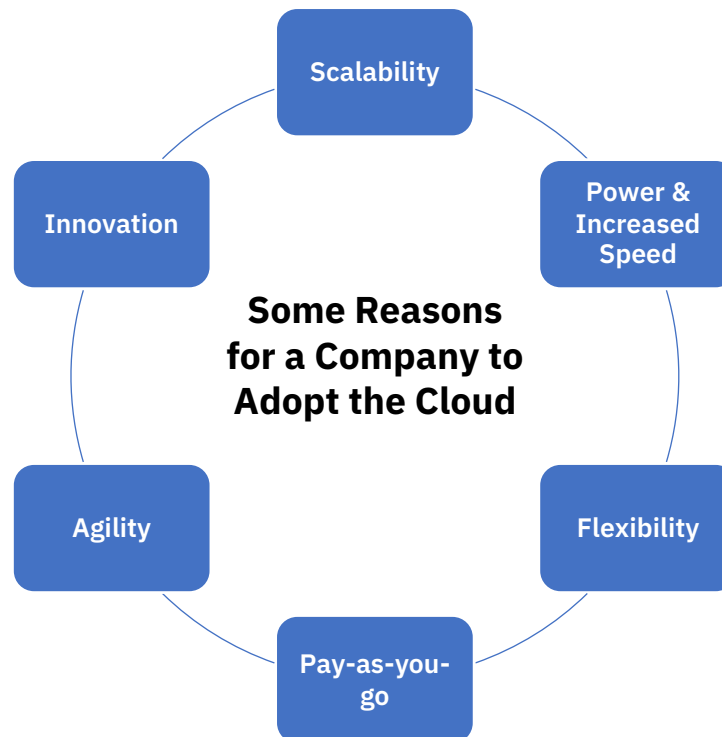- Identify where to find business case studies for cloud computing

Course 4: Business Case Study for Cloud Computing

Transformation and Modernization

BNP Paribas - Transformation

- Modern cloud concepts can assist an established bank in improving quality and efficiency and in saving money in the process.
- The core-systems modernization project included implementing IBM Developer for z/OS and IBM Connect for z/OS, which provided a way to bring the power of APIs to the mainframe.
  - BNP Paribas was able to utilize their existing Z applications in their new hybrid cloud strategy.
- These implementations changed the development workspace at BNP Paribas.
  - IT provides a modern environment with a standard, user-friendly interface.
  - Real-time code testing allows developers to catch and resolve potential code issues earlier.
- As BNP Paribas continues to modernize, they expect to see additional business value in terms of control, autonomy, quality, speed, and efficiency.

Scalability

Power & Increased Speed

Innovation

Some Reasons for a Company to Adopt the Cloud

Flexibility

Agility

Pay-as-you-go

Question 1.

A client would like to modernize their long-established core systems running on z/OS. How does the IBM Cloud Hybrid Cloud approach help this client?

A. Provides tools to automatically break up monolith applications into smaller microservices

B. Only cloud provider that allows clients to move production z/OS applications to the public cloud

C. Provides tools to automatically **migrate** applications to newer based platforms

D. Facilitates leveraging strategic core applications on legacy platforms while developing modern user-friendly interfaces in containers

Answer D. The IBM Cloud Hybrid approach facilitates leveraging strategic core applications on legacy platforms while developing modern user-friendly interfaces in containers.

Question 2.

A client is looking to retain their existing technology while moving towards more modern development practices and technologies. What benefit might IBM Cloud bring to this client?

A. Only cloud provider that allows clients to move production z/OS applications to the public cloud

B. Provides tools to automatically migrate applications to newer based platforms

C. Facilitates leveraging strategic core applications on legacy platforms while developing modern user-friendly interfaces in containers

D. Provides tools to automatically break up monolith applications into smaller microservices

Answer C. Facilitates leveraging strategic core applications on legacy platforms while developing modern user-friendly interfaces in containers

Question 3.

A client would like to modernize to improve the quality and responsiveness of their digital offerings. Their strategic core applications currently run on IBM Mainframe technology. What is the Hybrid Cloud approach for this modernization?

A. Modernize the entire application, using Watson AI from IBM Cloud and container and DevOps practices and tools from another cloud provider

B. Use IBM Cloud Wazi as a service to quickly modernize the legacy application and move it to Red Hat OpenShift on IBM Cloud

C. Move the entire application to the IBM Cloud z/OS LPARs as a service

D. Keep the mainframe application on-premises, while developing a user-friendly interface using cloud native methodologies

Answer D. Keep the mainframe application on-premises, while developing a user-friendly interface using cloud native methodologies

Question 4.

The attributes of a _____ cloud model include efficiency, autonomy, and modernization.

A.   Hybrid

B.   Private

C.   Public

D.   Concurrent

Answer A. The attributes of a hybrid cloud model include efficiency, autonomy, and modernization.

Acronyms

Acronym	Acronym Expansion
AI	Artificial Intelligence
BYOIP	Bring Your Own IP
CDN	Content Delivery Network
CIS	Center for Internet Security
DARPA	Defense Advanced Research Projects Agency
DISA	Defense Information Systems Agency
DoD	Department of Defense
DRAM	Dynamic Random Access Memory
FERPA	Family Educational Rights and Privacy Act
GDPR	General Protection Regulation
IaaS	Infrastructure as a Service
IP	Internet Protocol
KVM	Kernel-based Virtual Machine
MTCS	Multi-Tier Cloud Security
NIST	National Institute of Standards and Technology
PaaS	Platform as a Service
SaaS	Software as a Service
SRG	Security Requirement Guide
VM	Virtual Machine
VPC	Virtual Private Cloud
VSI	Virtual Server Instance

Courses and Objectives

Courses

1. Introduction to IBM Cloud
2. IBM Cloud Resources
3. IBM Cloud Services
4. Account Security
5. Cloud Services and DevSecOps
6. Cloud Database, Integration, and Analytics

Objectives

- Recognize the features and components of the IBM Cloud® platform
- Identify IBM Cloud resources available to support client solutions
- Match IBM Cloud Catalog services with their specific functionality
- Recall key security and compliance measures implemented in the IBM Cloud platform
- Identify elements of DevSecOps and how IBM Cloud services are delivered to teams
- Summarize core groups of available database, integration, and analytics services

OBJECTIVES

Course 1: Introduction to IBM Cloud

Introduction and Objectives

In Course 1 of the study guide, the subject matter:

- Examines the basic IBM Cloud® components, including the IBM Cloud platform, IBM Cloud platform security, and IBM Cloud hosting options.

Lessons

- IBM Cloud Platform
- IBM Cloud Platform Security
- IBM Cloud Hosting Options
- Summary
- Knowledge Check

Objective

- Recognize the various components of a cloud architecture and its relevance

Course 1: Introduction to IBM Cloud

The IBM Cloud Platform

The IBM Cloud platform combines the PaaS model with the IaaS model to provide clients with an integrated experience. This means that the software services and tools provided with PaaS are combined with computing services such as virtual servers, bare metal servers, networking, and storage. The IBM Cloud platform supports public, private, hybrid, multicloud, and virtual private cloud deployment models.

Key Features of the IBM Cloud Platform

- **Scalability**: The IBM Cloud platform can easily scale to support any client.

- **Reach and Reliability**: The IBM Cloud platform is deployed globally in data centers around the world.

- **Security**: IBM Cloud offers the most secure public cloud for business with a next-generation hybrid cloud platform, advanced data and AI capabilities, and deep enterprise expertise across 20 industries.

- **Open Source Technologies**: The IBM Cloud platform uses open source technologies, including Kubernetes and Red Hat® OpenShift® to provide clients with the control and flexibility needed to operate in a hybrid cloud environment.

- **Open Architecture**: The IBM Cloud platform is built on open architecture.

IBM Cloud Platform Architecture

Cloud platform services can be deployed across data centers around the world. Many components within the cloud platform work together to provide a consistent, dependable cloud experience for users.

Component	Definition
Console	Cloud resources can be created, viewed, and managed
Identity and Access Management	Authenticates users and controls access to resources
Catalog	Lists all the products available in the cloud
Search and Tagging	Filters and identifies resources (for example, keywords)
Account and Billing Management	Displays pricing plan usage and provides fraud protection

Course 1: Introduction to IBM Cloud

Data and Platform Security

The IBM Cloud platform was designed with security in mind. The platform itself provides layers of security control across the network and infrastructure.

IBM Cloud Compliance Programs and Certifications

IBM engages different firms to assess the security and compliance of the IBM Cloud platform. These detailed reviews and assessments from outside organizations allow the IBM Cloud platform to earn certifications for global, government, industry, and regional programs.

GDP Regulations

The General Data Protection Regulation (GDPR) is a regulation throughout the European Union that asserts an individual's right to privacy. IBM is working to embed data protection principles even more deeply into its business processes.

SOC Reports

Service Organization Control (SOC) reports demonstrate that IBM has appropriately designed its controls for the selected Trust Service Principles and that the controls operated effectively for the report period.

PSIRT

IBM Cloud also follows the IBM Product Security Incident Response Team (PSIRT) process for security incident management. This is a centralized process through which IBM clients, security researchers, industry groups, government organizations, or vendors report potential IBM security vulnerabilities.

Course 1: Introduction to IBM Cloud

IBM Cloud Data Centers

IBM offers a wide range of choices for when and how data and workloads should run. Workloads can be deployed in over 60 data centers, located in 19 countries, across six continents. This forms the cloud's highly resilient, available, and scalable foundation.

IBM Cloud Availability Zones and Regions

Logically isolated data centers with independent infrastructures are called zones. Regions are geographic areas that consist of zones. Regions provide a full cloud service stack enabling highly available, redundant, and geographically distributed client solutions. This provides the foundation for building and deploying mission critical applications in the cloud.

● Data Center & Network PoP   ◇ Multi-Zone Region   ● Network PoP   ○ Federal Data Center
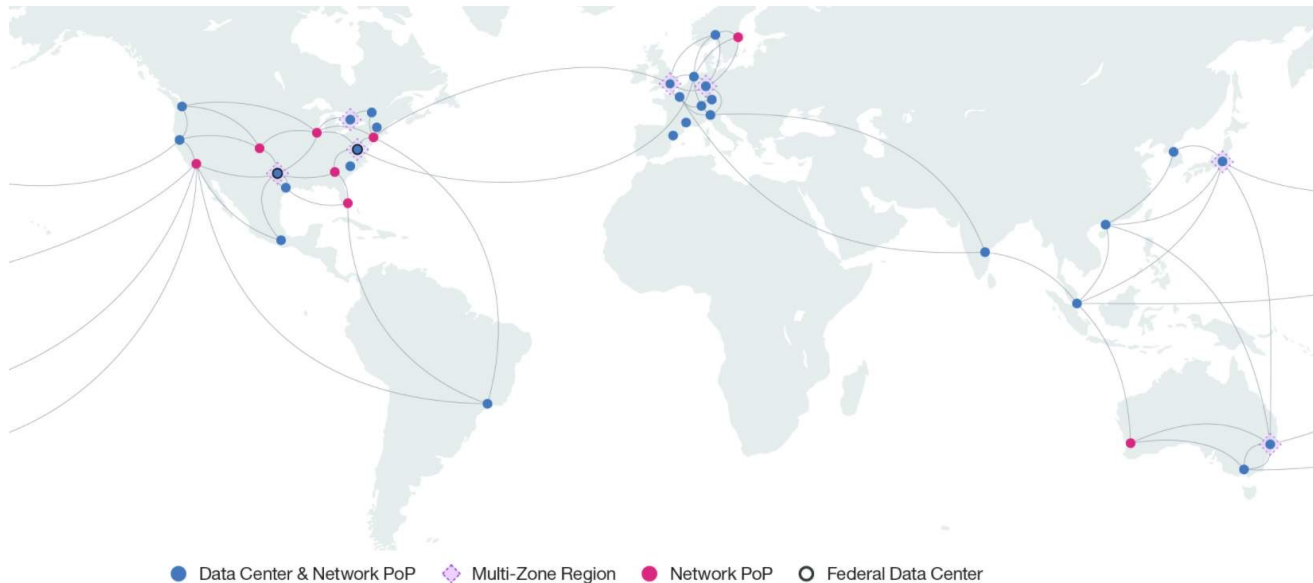
Resources are deployed in multiple zones to achieve fault tolerance and high availability. Because of this zone availability design, businesses can run critical workloads on the cloud with up to 99.999 percent availability.

Question 1.

The _____ is designed to enhance data protection.

A.   General Data Protection Regulation
B.   Generic Data Protection Regulation
C.   Global Data Protection Regulation
D.   Guarded Data Protection Regulation

Answer A. The General Data Protection Regulation is designed to enhance data protection.

Question 2.

The _____ ensures that IBM has designed the cloud controls as per the appropriate Trust Service Principles and is operating effectively for the reporting period.

A.   Security Reports

B.   Data Reports

C.   Service Organization Control Reports

D.   Compliance Reports

Answer C. The Service Organization Control Reports ensure that IBM has designed the cloud controls as per the appropriate Trust Service Principles and is operating effectively for the reporting period.

Question 3.

IBM Cloud follows which process for security incident management?

A.   Business Automation Security Incident Process

B.   ART Security Incident Process

C.   Product Security Incident Response Process

D.   Agile Security Incident Process

Answer C. IBM Cloud follows the Product Security Incident Response Process for security incident management.

Question 4.

In IBM Cloud, _____ provides a full cloud service stack that is highly available, redundant, and geographically distributed for building and deploying missing-critical client applications.

A.   Virtual private cloud

B.   Availability zones

C.   Regions

D.   Point of presence

Answer C. In IBM Cloud, regions provide a full cloud service stack that is highly available, redundant, and geographically distributed for building and deploying missing-critical client applications.

Question 5.

The IBM Cloud platform supports public, private, multicloud, virtual private, and what other cloud deployment model?

A.   Edge

B.   Hybrid

C.   Satellite

D.   Quantum

Answer B. The IBM Cloud platform supports public, private, multicloud, virtual private, and hybrid cloud deployment models.

Question 6.

An organization can achieve fault tolerance and high availability for their applications in IBM Cloud by doing what?

A.   By deploying resources in a single subnet

B.   By deploying resources in a single availability zone

C.   By deploying resources in a classic pod

D.   By deploying resources in multiple availability zones

Answer D. An organization can achieve fault tolerance and high availability for their applications in IBM Cloud by deploying resources in multiple availability zones.

Course 2: IBM Cloud Resources

Introduction and Objectives

In Course 2 of the study guide, the subject matter:

- Explains the IBM Cloud catalog, IBM Cloud documentation, and IBM Cloud tutorials.
- Reviews how to set up an IBM Cloud account and identify other resources that will assist in creating IBM Cloud solutions.
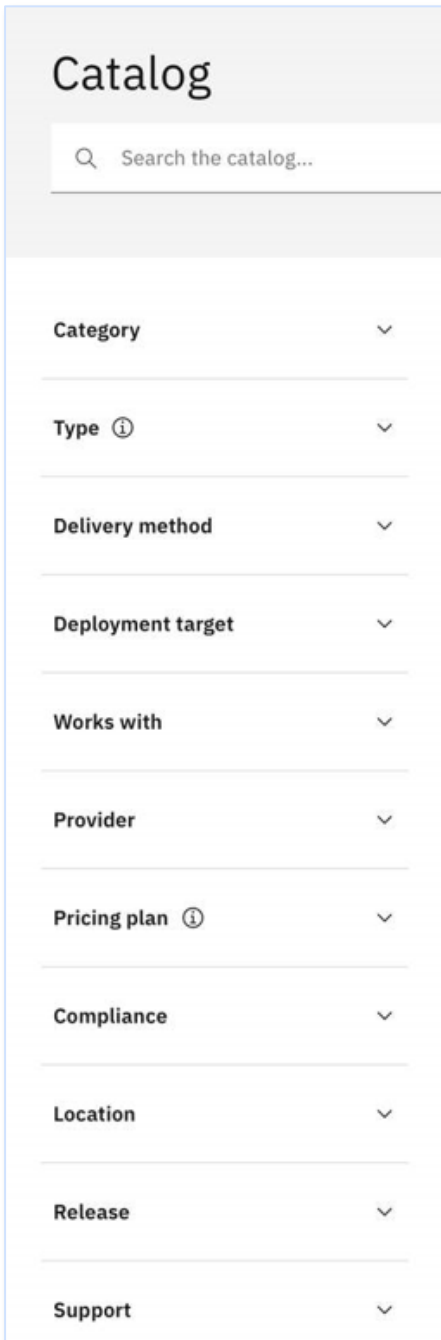
Lessons

- IBM Cloud Catalog
- IBM Cloud Documentation
- IBM Cloud Tutorials
- Additional IBM Cloud Resources
- Summary
- Knowledge Check

Objective

- Identify IBM Cloud resources available to support client solutions

Course 2: IBM Cloud Resources

IBM Cloud Catalog

There are a wide variety of IBM Cloud services available on IBM Cloud, both IBM-provided and created by third parties. They are found in the IBM Cloud catalog, which provides a central location for all services on IBM Cloud. To provision services from the catalog, users first search for the service, then click the service's tiles to start the provisioning process.

Users can filter services based on a variety of parameters including provider (IBM or third party), pricing plan (Lite or Free), and category (based on function).

The API Docs and general documentation can also be accessed using the appropriate links in the IBM Cloud catalog.

Once a user has selected a particular service, they can select the Docs link on the provisioning page to access the IBM Cloud documentation page for that service.

Course 2: IBM Cloud Resources

IBM Cloud Documentation

- The left-hand menu is used to navigate through major topics, such as tutorials, of a service.

- The right-hand menu in documentation is used to navigate the topics and sub-topics included in the currently opened documentation.

- Breadcrumbs at the top of the screen allow a user to navigate to a higher point within the IBM Cloud documentation navigation tree.

Course 2: IBM Cloud Resources

IBM Cloud Tutorials

All the services available in IBM Cloud are documented in the IBM Cloud Docs. If a deeper level information on any service in IBM Cloud is needed, the Docs pages are the best place to start. IBM Cloud Docs also includes a tutorial library for many of the services available in IBM Cloud. Search or filter by product, tutorial length, or services used to find what is needed.

Additional Sites

The following IBM websites also provide a wealth of information that help users design and develop applications using IBM Cloud services.

IBM Architecture Center

IBM's reference architecture is a source to help architect solutions that enable ways to meet business objectives using leading edge hybrid cloud and AI technologies.

IBM Developer

The IBM Developer site contains a collection of learning paths for step-by-step technical training, tutorials, articles, videos, and more.

Code Patterns

Code patterns are a collection of common use cases across multiple industries that involve multiple technologies, products, or services to solve issues that developer advocates have identified.

Question 1.

The IBM Cloud Catalog has many filters. What are two of them?

A. Storage

B. Deployment target

C. Location

D. Service Level

E. Country

Answer B and C. IBM Cloud Catalog has different filters including Category, Type, Delivery method, Location, Deployment target, and more.

Question 2.

A client is filtering for pricing plans while searching in the IBM Cloud Catalog. Which filters are available in the pricing plan?

A.   Enterprise and Personal plans

B.   Corporate plans

C.   Lite and Free plans

D.   Premium plans

Answer C. Lite and Free plans are the filters available in the pricing plans.

Question 3.

Which IBM Cloud resource provides best practices and industry examples to help a client looking to redesign and implement a hybrid cloud solution on IBM Cloud?

A.   Automation Center

B.   Architecture Center

C.   Security Center

D.   Migration Center

Answer B. The Architecture Center provides the best practices and industry examples to help a client looking to redesign and implement a hybrid cloud solution on IBM Cloud.

Question 4.

The IBM Developer site provides what kind of support?

A. A tool for reaching out directly to the development labs

B. A collection of learning paths, videos, and tutorials for the most demanded IBM products and services

C. A self-service application where a client can download fix packs directly from developers

D. An internal repository for IBM solution developers and programmers

Answer B. The IBM Developer site provides a collection of learning paths, videos, and tutorials for the most demanded IBM products and services.

Question 5.

The Code Patterns section of the IBM Developer site provides what type of material?

A.   A set of best practices for developers and system integrators

B.   A repository of widgets for low code programming

C.   A collection of articles with guides for senior developers

D.   A collection of common use cases with ready to use code

Answer D. The Code Patterns section of the IBM Developer site provides a collection of common use cases with ready to use code.

Course 3: IBM Cloud Services

Introduction and Objectives

In Course 3 of the study guide, the subject matter:

- Explains what types of services are included in the IBM Cloud catalog, including compute resources, container services, storage and database services, artificial intelligence and machine learning services, and security services.

Lessons

- IBM Cloud Catalog Services
- Compute Resources and Container Services
- Storage and Database Services
- Artificial Intelligence and Machine Learning Services
- Security Services
- Summary
- Knowledge Check

Objective

- Match IBM Cloud Catalog services with their specific functionality

Course 3: IBM Cloud Services

Compute Resources
Compute resources give clients access to servers and other resources that are needed to process information and run applications. Categories of compute resources include:

Bare Metal Servers
IBM Cloud Bare Metal Servers are dedicated, single tenant servers. These servers can be deployed in either a classic or virtual private cloud (VPC) infrastructure.

Virtual Servers
IBM Cloud Virtual Servers are scalable virtual servers that are purchased with cores and memory allocations. These servers are a great option for clients that need compute resources that can be added in minutes. These servers can be deployed in either a classic or virtual private cloud infrastructure.

IBM Power System Virtual Servers
IBM Power Systems Virtual Servers are a Power Systems offering. Power Systems Virtual Servers are located in IBM data centers but are separated from other IBM Cloud servers by separate networks and storage.

Serverless Architecture
Serverless is a cloud computing application development and execution model that enables developers to build and run application code without provisioning or managing servers and backend infrastructure. This doesn't mean there are not any servers, just that the servers are not visible to developers because the service provider (IBM) is managing the servers. IBM Cloud Code Engine is the name of IBM's serverless development environment.

Containers
Containers are packages of software that contain everything needed to run an application in any environment. They are a popular method of deploying applications to the cloud. IBM Cloud provides several services to assist organizations in managing and running container-based applications. These resources include IBM Kubernetes Service, Red Hat OpenShift on IBM Cloud, and Code Engine.

Course 3: IBM Cloud Services

Storage Services

The code in applications needs a place to store information, and IBM Cloud has several solutions including block storage, file storage, and IBM Cloud Object Storage. Backup and recovery options are also available.

Storage Type	Definition
Block Storage	Block storage is a technology that is used to store data files on Storage Area Networks (SANs) or cloud-based storage environments by breaking the data into blocks and storing those blocks as separate pieces.
File Storage	File storage is storage that is typically presented from Network Attached Storage (NAS) technology. Files are typically stored in a directory tree structure using folders, the same way they are viewed on a local hard drive.
IBM Cloud Object Storage	IBM Cloud Object Storage is a secure, reliable way to store data. Users have the option of storing data in a single data center, multiple data centers in a region, or across three regions for maximum safety. Data is stored in units called "buckets."
Backup and Recovery Options	There are several services available in IBM Cloud to help clients backup and quickly recover their most important workloads. Options include IBM Cloud Backup and third-party tools like Veeam and Zerto. IBM Spectrum Protect and some operating system specific options are available for IBM Power Systems Virtual Servers.

Course 3: IBM Cloud Services

Databases and Database Services

Storing structured data is often easier using a database rather than a traditional file system. IBM Cloud offers several databases, including two of the more popular databases:

- **Db2** is a relational database with enterprise-grade performance. It provides multizone support allowing it to be run in a highly available configuration.

- **Cloudant** is a document database that uses JSON files to store and index documents. Its serverless design allows for easy scaling based on need.

Artificial Intelligence (AI) and Machine Learning (ML) Services

IBM Watson® is IBM's artificial intelligence platform. Clients can use AI and ML to enhance user interactions, search functionality, and content analytics in their applications. Some of the most common services are:

- **Watson Assistant** can add a conversational chat functionality to an application.

- **Watson Discovery** is used to gain a deeper understanding of data and implement natural language document search capabilities.

- **Watson Studio** allows customization of machine learning models to optimize decisions using open source tools.

- **Watson OpenScale** is designed to track and measure the performance of machine learning models.

Security Services

With a wide variety of bad actors putting data at risk, security is a high priority. IBM Cloud provides many solutions to protect data. Security features in IBM Cloud include:

- **IBM Key Protect** for IBM Cloud is a service for managing cryptographic keys, which are used to protect data.

- **Hyper Protect Crypto Services** provides dedicated services and hardware security modules (HSMs) to manage client crypto keys.

- **Hyper Protect Servers** allow clients to create and run virtual servers on IBM LinuxONE®, the industry's most secure Linux-based platform.

- **Firewalls and VPC Solutions** assist clients in blocking unauthorized access to both public and private clouds, while still permitting outward communications.

Question 1.

A client is interested in an application but doesn't want to be concerned with provisioning any infrastructure resources. Which IBM Cloud service meets this client's requirement?

A.   Code Engine

B.   Red Hat OpenShift on IBM Cloud

C.   Kubernetes Service

D.   IBM Power Virtual Servers

Answer A. IBM Cloud Code Engine is the name of IBM's serverless development environment.

Question 2.

Red Hat OpenShift on IBM Cloud and _____ assist organizations in managing and running container-based applications.

A. VMware Solutions

B. Virtual Server for VPC

C. Analytics Engine

D. Code Engine

Answer D. Red Hat OpenShift on IBM Cloud and Code Engine assist organizations in managing and running container-based applications.

Question 3.

_____ stores data in buckets.

A.   Spectrum Protect

B.   Block Storage

C.   Object Storage

D.   File Storage

Answer C. Object Storage stores data in buckets.

Question 4.

A client needs a service that will regularly back up its customer's data so that it can easily be recovered in case of data corruption. Which service should be recommended to this client?

A.   Block Storage

B.   Schematics

C.   Cloud HSM

D.   Spectrum Protect

Answer D. Spectrum Protect will help a client's need to regularly back up its customer's data so that it can easily be recovered in case of data corruption.

Question 5.

_____ provides a document database that uses JSON files to store and index documents.

A. Databases for MySQL

B. Cloudant

C. Db2

D. DataStage

Answer B. Cloudant provides a document database that uses JSON files to store and index documents.

Question 6.

Which of the following can be used to track and measure the performance of machine models that can safeguard data models against biases?

A.   Watson OpenScale

B.   Watson Studio

C.   Watson Assistant

D.   Watson Discovery

Answer A. Watson OpenScale can be used to track and measure the performance of machine models that can safeguard data models against biases.

Question 7.

_____ allows natural language documentation search.

A.   Watson OpenScale

B.   Watson Search

C.   Watson Assistant

D.   Watson Discovery

Answer D. Watson Discovery allows natural language documentation search.

Question 8.

_____ provides dedicated services and hardware security modules to manage client crypto keys.

A.   Key Protect

B.   Privileged Access Gateway

C.   Hyper Protect Crypto Services

D.   Secrets Manager

Answer C. Hyper Protect Crypto Services provides dedicated services and hardware security modules to manage client crypto keys.

Course 4: Account Security

Introduction and Objectives

In Course 4 of the study guide, the subject matter:

- Explains the threats associated with cloud-based solutions and explores IBM Cloud security features and other network protocols that can be used to mitigate these threats.

Lessons

- IBM Cloud Security Strategy Overview
- Shared Responsibilities
- Identity and Access Management
- Encryption
- Regulatory and Compliance Programs
- Summary
- Knowledge Check

Objective

- Recall key security and compliance measures implemented in the IBM Cloud platform

Course 4: Account Security

Cloud-Based Security Threats

Since cloud-based solutions reside in a data center that is owned by another company, ensuring security in cloud-based applications is critical. A few potential threats organizations may face are data breaches, malware injections, Distributed Denial of Service (DDoS), Advanced Persistent Threats (APTs), insecure API, and account hijacking.

If data is breached, it may cause fines and penalties from regulatory agencies. In some situations, these can be steep. The cloud shared responsibility model defines who is responsible for what. The cloud provider is responsible for security of the cloud environment, and the client is responsible for security of the services and applications they have on the cloud.

Identity and Access Management

Identity and Access Management (IAM) is the process of ensuring the right resources have access to the right assets (tools, applications, data, machines, and so on) at the right time and for the right reasons. Proper identity and access management processes allow resources to work productively and at the same time allow organizations to protect their data and other assets.

IAM allows administrators to assign user-specific access rights. In addition, administrators may grant privileges to users on an individual basis, or they may grant access to large numbers of users (access group) at the same time. Access groups are assigned to an access policy, which determines the role of the user. Roles are mapped to actions each role is allowed to perform, for example, adding, removing, or deleting services and instances.

Course 4: Account Security

Data Encryption

Encryption is the process of encoding data, so it is hidden and inaccessible to unauthorized users. Encrypted data may only be accessed and read using keys. Clients own the keys necessary to access their data. This allows clients to ensure that their data is secure and that they are the only ones allowed access to it. The following are the three encryption options IBM Cloud offers to ensure the safety and security of data.

Data at rest is data that is being stored for future use. IBM Cloud Storage services are used to protect data at rest.

Data in motion is data that is in the process of being transferred. IBM Cloud Networking services can be used to protect data in motion.

Data in use is data that is currently being processed. IBM Cloud Compute services can be used to protect data in this state.

Course 4: Account Security

IBM Cloud Compliance

In addition to ensuring applications and data are safe from external and internal threats, cloud solutions must comply with any government, regional, or industry regulations imposed on the organization. The compliance programs are divided into four categories: global compliance, government compliance, industry compliance, and regional compliance.

Global Compliance

IBM Cloud provides programs and certifications that help organizations establish compliance for a variety of international standards. Many standards are included in the Center for Internet Security (CIS) Benchmarks. The CIS IBM Cloud Foundations Benchmark provides organizations with secure configuration guidelines for adopting IBM Cloud services. It's a collection of industry best practices for securely configuring IT systems, software, and networks.

Government Compliance

IBM Cloud offers two platforms designed specifically for government and public sector clients:
- IBM Cloud for Government (IC4G)
- IBM SmartCloud® for Government (SCG)

Industry Compliance

IBM Cloud infrastructure and Platform-as-a-Service (PaaS) offerings help organizations meet sector-specific compliance requirements by providing services that support key industry programs.

Regional Compliance

International organizations must adhere to region specific compliance standards when they move their IT infrastructure to the cloud.

Question 1.

Which process can an organization use to ensure the right resources have the access to the right assets?

A.   Incident and Operations Management

B.   Identity and Access Management (IAM)

C.   Change Management

D.   Security and Regulatory Compliance

Answer B. IAM is the process that an organization can use to ensure the right resources have access to the right assets.

Question 2.

When using IBM Cloud, _____ encryption protects data that is being stored for future use.

    A.   Data in motion

    B.   Data at rest

    C.   Data in stasis

    D.   Data in use

Answer B. When using IBM Cloud, data at rest encryption protects data that is being stored for future use.

Question 3.

_____ encryption protects data on IBM Cloud that is in the process of being transferred.

A. Data in use

B. Data in stasis

C. Data in motion

D. Data at rest

Answer C. Data in motion encryption protects data on IBM Cloud that is in the process of being transferred.

Question 4.

_____ encryption protects data on IBM Cloud that is currently being processed.

A.   Data in stasis
B.   Data in motion
C.   Data at rest
D.   Data in use

Answer D. Data in use encryption protects data on IBM Cloud that is currently being processed.

Question 5.

_____ services can be used to protect data at rest.

A.   Storage
B.   Observability
C.   Networking
D.   Compute

Answer A. Storage services can be used to protect data at rest.

Question 6.

_____ services can be used to protect data in motion.

A. Storage
B. Observability
C. Networking
D. Compute

Answer C. Networking services can be used to protect data in motion.

Question 7.

_____ services can be used to protect data in use.

A.   Storage

B.   Observability

C.   Networking

D.   Compute

Answer D. Compute services can be used to protect data in use.

Question 8.

A collection of industry best practices for securely configuring IT systems, software, and networks is called what?

A. Department of Defense (DoD) Cloud Computing Security Requirement Guide (SRG)

B. US Health Insurance Portability and Accountability Act (HIPAA)

C. The Family Educational Rights and Privacy Act (FERPA)

D. Center for Internet Security (CIS) Benchmarks

Answer D. The CIS Benchmarks is a collection of industry best practices for securely configuring IT systems, software, and networks.

Question 9.

Which of the following security features will allow an IBM Cloud account administrator to limit who can delete a service instance?

A.   Center for Internet Security (CIS) IBM Cloud Foundations Benchmark

B.   Identity and Access Management (IAM)

C.   Transport Layer Security (TLS)

D.   Federal Risk and Authorization Management Program (FedRAMP)

Answer B. IAM is the security feature that will allow an IBM Cloud account administrator to limit who can delete a service instance.

Question 10.

Who is responsible for application data in the IBM Cloud shared responsibility model?

A. Client

B. IBM

C. Client and IBM

D. Individual user

Answer A. The client is responsible for application data in IBM Cloud shared responsibility model.

Course 5: Cloud Services and DevSecOps

Introduction and Objectives

In Course 5 of the study guide, the subject matter:

- Reviews and contrasts the differences between DevOps and DevSecOps.
- Explains Continuous Delivery.

Lessons

- DevOps Versus DevSecOps
- Continuous Delivery
- Summary
- Knowledge Check

Objective

- Identify elements of DevSecOps and how IBM Cloud services are delivered to teams

Course 5: Cloud Services and DevSecOps

Background

In the past, developers would develop an application then pass it along to the IT operations team. This usually led to many change requests after the application had been released and to make it work correctly in the target environment.

DevOps

DevOps is formed from the words "Development" and "Operations." The main goal of this new methodology was to shorten the development cycle of applications and increase their quality by incorporating feedback from IT Operations in the development process, rather than getting operations feedback once the application was completed. DevOps outlines an organizational culture shift and development process that helps in meeting a user's demand for new features and uninterrupted availability and performance.

Development → ● ← Operations
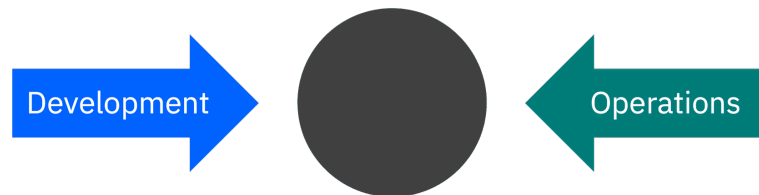
DevSecOps

Since running applications over the internet became widespread, the need to integrate security within DevOps became more important. This created a methodology known as DevSecOps. In this methodology, security is integrated into every phase of development.

Security

Development → ● ← Operations
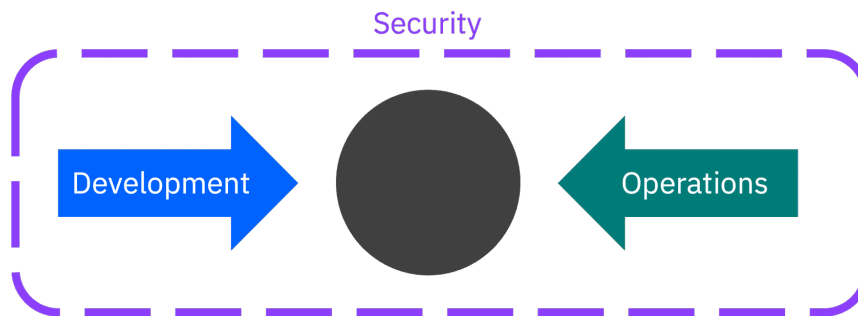
DevOps and DevSecOps was developed with a continuous development cycle in mind. Instead of developing an application then releasing an update a couple times a year, DevOps is designed to continually monitor and improve applications.

Course 5: Cloud Services and DevSecOps

Phases of the DevSecOps Cycle
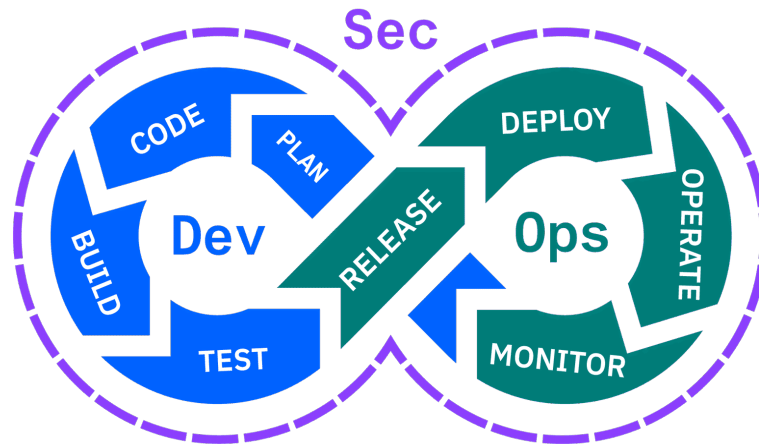
Continuous Delivery is a service that is used to speed up the process of delivering development changes to production by providing tools to automate the building and deployment of applications. It can be used to deploy a Kubernetes application on IBM Cloud.

Toolchains

- Set of tools used to automate the process of building, deploying, and managing applications

- Can create toolchains from a variety of sources including other IBM Cloud services, third-party software, and open source tools

Automated Pipelines

- Can be set up to automate the process of building, testing, and deploying software

- Allows for updates to existing software with minimal human intervention

Code Anywhere

- Uses an integrated web-based environment

- Can create, edit, run, and debug code from anywhere in the world

Question 2.

Why would an organization use a toolchain when deploying on IBM Cloud?

- A. To link user applications so that data can easily flow between them programmatically
- B. To record transaction in an immutable ledger that is highly transparent and secure
- C. To automate the process of building, deploying, and managing the application life cycle
- D. To enable cloud administrators to group sets of users together and grant them common privileges

Answer C. An organization would use a toolchain when deploying IBM Cloud to automate the process of building, deploying, and managing the application life cycle.

Question 3.

A client could use the _____ methodology in IBM Cloud to speed and secure development.

A.   Terraform

B.   DevSecOps

C.   Jenkins

D.   IBM Agile Method

Answer B. A client could use the DevSecOps methodology in IBM Cloud to speed and secure development.

Course 6: Cloud Database, Integration, and Analytic Services

Introduction and Objectives

In Course 6 of the study guide, the subject matter:

- Explores the databases that are available on IBM Cloud.
- Explores the services available for analytics and the services that are used to connect apps together.

Lessons

- Databases on IBM Cloud
- Analytics of IBM Cloud
- Integration on IBM Cloud
- Summary
- Knowledge Check

Objective

- Summarize core groups of available database, integration, and analytics services

Course 6: Cloud Database, Integration, and Analytic Services

Types of Databases
IBM provides four types of databases.

Relational databases are best at storing data that can be organized into a relational table structure. Examples include: DB2®, PostgreSQL, EnterpriseDB (EDB) Postgres, and MySQL.

Document databases are best at storing semi-structured data like user-generated content and order data. Examples include MongoDB, Elasticsearch, and Redis.

Key value databases are best at storing data that is structured but doesn't require the relational abilities of a relational database. Data are stored as a key value pair where the key is a unique identifier. An example includes etcd.

Columnar databases store data in columns instead of rows. They are helpful for data analytics and data warehousing as they speed up the time it takes to return a query. DataStax is an example of a columnar database.

IBM Cloud Database Use Cases

SAP, ERPs, CRMS, and ledgers	**Databases:** **PostgreSQL** - customizable, open source object-relational databases **EnterpriseDB** - optimizes the built-in features of PostgreSQL; adds compatibility with Oracle **Db2® on Cloud** - based on IBM Db2 database engine; provides a fully managed solution
Web and Mobile Apps	**IBM Cloudant®** - a scalable, document database that uses JSON **MongoDB** - JSON-based document-oriented database, which includes rich query functionality **DataStax** - NoSQL database built on Apache Cassandra
Developer Tools	**Elastisearch** - based on JSON document databases and allows full-text search **Redis** - designed for in-memory functionality making them very fast
Confidential Data	**Hyper Protect** - can be used with PostgreSQL and MongoDB for fully managed, highly secure applications
Business Intelligence	**IBM Cloud Databases for PostgreSQL** - database warehouse optimized for machine learning and high-performance analytics

Course 6: Cloud Database, Integration, and Analytic Services

IBM Cloud Analytics are designed to work together with other products to provide new and unexpected insights into data in order to deliver business-changing results.

There are three types of analytics:
- **Descriptive** – method of collecting data from a source to identify patterns or trends
- **Diagnostic** – further analyzes data with a goal of identifying the root of any patterns
- **Prescriptive** – applies data models with a goal of projecting the data out into the future

Analytics Available on IBM Cloud

IBM Analytics Engine
- This builds on functionality of the open source technologies, Apache Hadoop and Apache Spark, but separates compute and storage functionality.
- Key Features: Open source, scales on demand, easy configuration
- Use cases: Log analytics and customer insights

IBM InfoSphere Information Server on Cloud
- The goal is to provide one platform that will integrate and analyze a wide variety of data.
- Key Features: Total data integration, standardized data, easily scalable
- Use cases: Data integration, Data quality

Course 6: Cloud Database, Integration, and Analytic Services

Integration Services Available in IBM Cloud

API Connect®

- Used to create and manage APIs in applications
- Integration tool that assists clients to package their data
- Key Features: intuitive development tools, easy yet powerful security, life cycle management, customizable developer portal

App Connect

- Used to connect various applications to each other
- Key Features: workflow automation, easy integration, create API flows without code

Event Streams

- Built on Apache Kafka and are used as high throughput message buses
- Key Features: fully managed, intuitive UX, security and compliance

MQ

- Provides enterprise-grade messaging capabilities between applications
- Key Features: deduplication (elimination of duplicate or redundant information), connectivity, encryption

Question 1.

MongoDB is considered a document database. What is another example of a document database?

    A.   EnterpriseDB Postgres

    B.   IBM Event Streams

    C.   etcd

    D.   Elasticsearch

Answer D. MongoDB and Elasticsearch are considered document databases.

Question 2.

The _____ could be used to store JSON structured data for an online shopping application.

A.   IBM Cloud Block Storage

B.   IBM Cloud etcd

C.   IBM Cloud Databases for MongoDB

D.   IBM Cloud Databases for MySQL

Answer C. The IBM Cloud Databases for MongoDB could be used to store JSON structured data for an online shopping application.

Question 3.

On IBM Cloud, _____ could be used to manage the API's life cycle and assist in designing the APIs and their rules.

A.   IBM Event Streams

B.   IBM DataPower API Gateway

C.   IBM API Connect

D.   IBM App Connect

Answer C. On IBM Cloud, IBM API Connect could be used to manage the API's life cycle and assist in designing the APIs and their rules.

Question 4.

On IBM Cloud, _____ provides a messaging system between multiple systems, could be in a different cluster and zones, handles connectivity issues, and encryption of end-to-end data.

A.   IBM Aspera

B.   IBM Object Storage

C.   IBM API Connect

D.   IBM MQ

Answer D. On IBM Cloud, IBM MQ provides a messaging system between multiple systems, could be in a different cluster and zones, handles connectivity issues, and encryption of end-to-end data.

Question 5.

The _____ service is built on Apache Kafka and provides a durable message storage and allows the interested parties to publish and subscribe to its topics.

A. IBM MQ

B. IBM Event Streams

C. IBM API Connect

D. IBM App Connect

Acronyms

Acronym	Acronym Expansion
AI	Artificial Intelligence
API	Application Programming Interface
APT	Advanced Persistent Threats
CSA	Cloud Security Alliance
CSF	Cybersecurity Framework
DBaaS	Databases as a Service
DDoS	Distributed Denial-of-Service
DevOps	Development, Operations
DevSecOps	Development, Security, Operations
DISA	Defense Information Systems Agency
DoD	Department of Defense
DSS	Data Security Standard
EDB	EnterpriseDB
FedRAMP	Federal Risk and Authorization Management Program
FERPA	Family Educational Rights and Privacy
FFIEC	Federal Financial Institutions Examination Council
FISC	Financial Industry Information Systems
FIPS	Federal Information Processing Standards
FISMA	Federal Information Security Management Act
HIPAA	Health Insurance Portability Accountability Act
HSM	Hardware Security Module

Acronyms

Acronym	Acronym Expansion
IaaS	Infrastructure as a Service
IAM	Identity and Access Management
IOS	International Organization of Standardization
IoT	Internet of Things
ITAR	International Traffic in Arms Regulations
JSON	JavaScript Object Notation
ML	Machine Learning
MTCS	Multi-Tier Cloud Security
NIST	National Institute of Standards and Technology
PaaS	Platform as a Service
PCI	Payment Card Industry
SaaS	Software as a Service
SANs	Storage Area Networks
SCG	SmartCloud for Government
SEC	US Security and Exchange Commission
SOC	Service Organization Control
SRG	Security Requirement Guide
SSL	Secure Sockets Layer
TLS	Transport Layer Security
UAT	User Acceptance Testing
VPC	Virtual Private Cloud

Courses and Objectives

Courses

1. Cloud Networking Components
2. Cloud Networks
3. Advanced Networking

Objectives

- Identify IBM Cloud® networking components and list their advantages
- Distinguish between classic and VPC networks
- Identify IBM Cloud network security services and options

Course 1: Cloud Networking Components

Introduction and Objectives

In Course 1 of the study guide, the subject matter:

- Introduces basic cloud networking components and explains how those components work together to form secure and robust cloud networks.

Lessons

- Network Overview
- Summary
- Knowledge Check

Objective

- Identify IBM Cloud networking components and list their advantages

Course 1: Cloud Networking Components

Network Types

A computer network is made of two or more computers that are connected either by cables (wired) or WiFi (wireless). A network's purpose is to transmit, exchange, or share data and resources. It is built by using hardware (for example, routers, switches, access points, and cables) and software (for example, operating systems or business applications).

Common Network Types

Local Area Network (LAN)
LANs connect computers over a relatively short distance. For example, a LAN may connect all the computers in an office building, school, or hospital. LANs are generally privately owned and managed.

Wireless Local Area Network (WLAN)
WLANs connect computers over a relatively short distance but do so wirelessly.

Wide Area Network (WAN)
WANs connect computers over a wide area, such as from region to region, or continent to continent. The internet is an example of a WAN. Typically, there is a collective or distributed ownership of WANs.

Virtual Private Network (VPN)
VPNs are a secure, point-to-point connection between two network end points. VPNs establish encrypted channels that keep users' identity and access credentials as well as any data transferred inaccessible to hackers.

Course 1: Cloud Networking Components

Network Terminology
Review the important terms and concepts below.

Term	Definition
IP Address	The Internet Protocol (IP) address is an identifier assigned to every device connected to a network that uses IP for communication.
Ports and Protocols	A port identifies a specific connection between network devices. Applications and services use port numbers to determine which application, service, or process should receive specific messages.
Nodes	A node is a connection point inside a network that can receive, send, create, or store data. Each node requires you to provide some form of identification to receive access, like an IP address.
Switch	A switch is a device that connects other devices and manages node-to-node communication within a network, ensuring data packets reach their ultimate destination.
Router	A router is a physical or virtual device that sends information contained in data packets between networks.
Gateway	A gateway is a piece of networking hardware or software that allows data to flow from one discrete network to another.
Firewall	A firewall is a network security device that monitors traffic to or from the network. Traffic is allowed or blocked based on a defined set of rules.
Load Balancer	A load balancer is a device that distributes network or application traffic across several servers.

Course 1: Cloud Networking Components

Network Security
Threats can enter the network through hardware that makes up the network, software that is used on the network, and devices that are used to access the network, such as computers, smartphones, or tablets.

Basic Security Measures
Basic security measures include:
- Using user IDs and passwords, which make proprietary or personal information more difficult to access.
- Installing updates and patches to software, which ensures the software environment is up to date and includes the latest bug fixes and security enhancements.
- Educating users on their role and responsibility in keeping the network secure, which ensures everyone is aligned with the tasks necessary to secure the network.
- Remaining aware of external threats, which ensures all users are on guard and watching for potential hacks.

Protection of Data in Transit
Data that is moving from one place to another, such as when it is transmitted over the network is referred to as data in transit. There are two main methods to protect data in transit:

Internet Protocols

When using IBM Cloud, organizations may employ internet protocols such as TLS (Transport Layer Security) and HTTPS (hypertext transfer protocol secure) to protect data as it transmitted.

IBM Cloud Services

Organizations may also provision IBM Cloud products and services that assist them in creating a more secure environment in which to transport data. Devices such as physical or virtual firewalls and routers might be paired with gateway appliances, VPNs, and services such as IBM Cloud Direct Link to create secure tunnels through which data can pass from one environment or infrastructure to another.

Question 1.

What is the definition of a virtual private network (VPN)?

A. VPNs are an unencrypted point-to-point connection between two network end points that establish a route used to transfer data behind a firewall.

B. VPNs are an unencrypted point-to-point connection between two clouds that establish a channel allowing data to be transferred enabling multi-cloud communications.

C. VPNs are secure, multi-point connections between three or more networks to establish encrypted channels.

D. VPNs are a secure, point-to-point connection between two network end points that establish encrypted channels.

Answer D. VPNs are a secure, point-to-point connection between two network end points that establish encrypted channels.

Question 2.

What is the definition of a wide area network (WAN)?

A.   A point-to-point encrypted connection allowing computers in different clouds to communicate

B.   Connects computers over a wide area using only point-to-point microwave or cellular communications

C.   Connects computers over a wide area, such as from region to region, or continent to continent

D.   Connects computers over campus sized areas, such as universities and large corporations

Answer C. A WAN connects computers over a wide area, such as from region to region, or continent to continent.

Question 3.

An IP address is defined as _____ .

A. A unique identifier assigned to every public device connected to a network that uses SSL/TLS for communication

B. A unique identifier assigned to every device connected to a network that uses IP for communication

C. A packet size specification assigned to every device connected to a network that uses IP for communication

D. A unique identifier assigned to every device connected to a network and is used to encrypt data sent to that device

Answer B. An IP address is a unique identifier assigned to every device connected to a network that uses IP for communication.

Question 4.

Describe how a port is utilized for communication.

A.   Applications or services use port numbers to receive specific messages

B.   A container destination that determines the location where messages are sent and received

C.   An application that determines which computer service or process should receive specific messages

D.   A specific router connection used to determine which port should send and receive specific messages

Answer A. Applications or services use port numbers to receive specific messages.

Question 5.

What is a network router used for?

A. A device that connects devices node-to-node within a network

B. A device that sends information between networks

C. A device that sends information between cloud networks

D. A router that connects devices node-to-node within a network and sends information between networks

Answer B. A network router's purpose is to send information between networks.

Question 6.

_____ is the only way to secure a connection to the applications deployed on IBM Cloud via a web server.

A.   FTP

B.   SSH

C.   TLS

D.   TELNET

Answer C. TLS is the only way to secure a connection to the applications deployed on IBM Cloud via a web server.

Course 2: Cloud Networks

Introduction and Objectives

In Course 2 of the study guide, the subject matter:

- Introduces IBM Cloud Classic and Virtual Private Cloud (VPC) networking components and their advantages.

Lessons

- IBM Cloud Classic Network Infrastructure
- IBM Cloud Virtual Private Cloud Network Infrastructure
- Summary
- Knowledge Check

Objective

- Distinguish between Classic and VPC networks

Course 2: Cloud Networks

IBM Cloud Classic Networking

The architecture, security, and connectivity of IBM Cloud Classic Networking is reviewed below.

Architecture

Network architecture includes both the physical and logical framework of a computer network. In IBM Cloud Classic networking, the physical framework is the data center. The logical framework is created using Virtual Local Area Networks (VLANs) and subnets

Security

Firewalls and gateways are two commonly implemented security options for IBM Cloud Classic networks.

- **Firewalls** are network security devices that monitor network traffic and prevent access to parts of the network based on security rules.
- **Gateways** can by physical or virtual devices. Associating a VLAN to a gateway reroutes traffic on that VLAN and all its subnets to the gateway. Traffic attempting to enter the VLAN must also go through the gateway.

In IBM Cloud Classic infrastructure, IBM Cloud provides both firewall and gateway functionality in the same device. Clients can use a Vyatta, Juniper vSRX, or FortiGate device. They can also bring their own device and deploy it on IBM Cloud hardware through an IBM Cloud bring-your-own-gateway service.

Course 2: Cloud Networks

IBM Cloud Classic Networking (Continued)

Connectivity

The IBM Cloud Classic network is able to provide a global footprint due to the three distinct and redundant network architectures it is built on. The three architectures include:

- A **public network** is used to connect infrastructure within the IBM Cloud Classic network to the internet.
- A **private network** consists of the IBM Cloud global private network, also called the IBM Cloud backbone.
- A **management network** is available for administrative purposes.

To allow a client to connect their IBM Cloud network to on-premises data centers or to other cloud service providers, IBM Cloud offers services such as IBM Cloud Direct Link and virtual private network (VPN) tunnels.

IBM Cloud Direct Link connects an organization's IBM Cloud network to on-premises and other cloud resources with consistent, higher-throughput connectivity. This is done without routing through the public internet.

A **VPN** establishes a secure connection between a client's network and IBM Cloud. This creates a secure connection between two networks using the public internet.

Typically, clients set up both IBM Cloud Direct Link and VPN connectivity between their environment and IBM Cloud. IBM Cloud Direct Link is generally used for all production traffic due to the consistent throughput and privacy, and the VPN is available just in case something causes the IBM Cloud Direct Link connection to go down.

Course 2: Cloud Networks

IBM Cloud Virtual Private Cloud (VPC) Infrastructure

The architecture, security, and connectivity of IBM Cloud VPC is reviewed below.

As a refresher, a VPC provides a private cloud environment on a shared public cloud infrastructure, where each user's data and workloads are separated from that of other users. The VPC private backbone is used to communicate between two virtual server instances within a VPC.

VPC

Architecture

Unlike IBM Cloud Classic architecture, the location of the VPC is not defined by data centers. Instead, the location is defined by zones and regions.

- A **region** is the geographic area in which a VPC is deployed.

- **Zones** are physical data centers. Data centers include the power, cooling, compute, networking, and storage resources used by services and applications. Each zone is isolated from the other zones.

    - **Subnets** are specified IP address range blocks, also known as Classless Inter-Domain Routing (CIDR) blocks. Subnets are bound to a single zone.

Security

IBM Cloud VPC workloads must be protected and secured. Security groups and network access control lists (ACLs) are two ways of controlling network traffic and limited access to resources inside of a VPC.

- **Access Control Lists (ACLs)** manage both inbound and outbound traffic for a subnet. This means ACLs allow or deny access to a subnet. Access is granted or denied based on a set of rules created by an administrator. Separate rules are allowed for inbound and outbound traffic.

- A **security group** acts as a virtual firewall and controls network traffic for one or more virtual server instances (VSIs). A security group is a collection of rules that specify what traffic to allow for an associated instance.

Course 2: Cloud Networks

IBM Cloud Virtual Private Cloud (VPC) Infrastructure (Continued)

Connectivity

Options for connecting resources on a subnet with resources on the internet include:

- **Public Gateway** enables a subnet and any resources attached to the subnet to connect to the internet. The gateway provides only a **one-way** connection. It does not allow traffic from the internet to connect with a resource on the subnet.
- **Floating IP addresses** are IP addresses that are provided by IBM Cloud. They can initiate or receive connections **to** or **from** the internet.

Cloud Connectivity

Most organizations will have multiple VPCs in different regions, multiple VPCs in the same region, or they may need to connect a VPC with resources in the Classic infrastructure network.

Using IBM Cloud **Transit Gateway**, organizations can create single or multiple transit gateways that connect specific VPCs together. They can also connect their IBM Cloud Classic infrastructure to a transit gateway to provide seamless communication with Classic infrastructure resources. Two methods of connecting resources include:

- **IBM Cloud Direct Link** solution is designed to seamlessly connect on-premises resources to cloud resources.
- Virtual Private Network **(VPN) gateways** connect two or more devices or networks together in a VPN infrastructure. These devices do not have to be located in the same VPC or even the same region.

Question 1.

Juniper vSRX is one gateway device that can be used on the IBM Cloud Classic network to safely transfer data between servers in different locations. What is another gateway device?

A. IBM Cloud Virtual Router

B. IBM Cloud Public Gateway

C. FortiGate Security Appliance

D. IBM Cloud Direct Link

Answer C. FortiGate Security Appliance is another gateway device that can be used on the IBM Cloud Classic network to safely transfer data between servers in different locations.

Question 2.

Which IBM Cloud offering will allow a company to have a private, secure, high-throughput connection to IBM Cloud?

A.   IBM Cloud Public Gateway

B.   IBM Cloud Transit Gateway

C.   IBM Cloud Direct Link

D.   IBM Cloud VPN

Answer C. The IBM Cloud Direct Link offering will allow a company to have a private, secure, high-throughput connection to IBM Cloud.

Question 3.

The geographic area in which an IBM Cloud VPC is deployed is called a _____ .

A. Region
B. Locale
C. Data Center
D. Zone

Answer A. The geographic area in which an IBM Cloud VPC is deployed is called a region.

Question 4.

In IBM Cloud VPC, a security group acts as a _____ .

    A.   Virtual private endpoint

    B.   Gateway

    C.   Firewall

    D.   Router

Answer C. In IBM Cloud VPC, a security group acts as a firewall.

Question 5.

If a client needs to securely connect to an IBM Cloud VPC over the internet, what service should they use?

A.   Direct Link

B.   Transit Gateway

C.   Public Gateway

D.   VPN Gateway

Answer D. A client should use VPN Gateway if they need to securely connect to an IBM Cloud VPC over the internet.

Question 6.

If a company deployed five virtual server instances in an IBM Cloud VPC, and these instances require outbound connectivity to an external service over the internet, which IBM Cloud VPC offering would be the best solution?

A.   Direct Link

B.   Transit Gateway

C.   Public Gateway

D.   Floating IP

Answer C. The company should use a Public Gateway.

Question 7.

In order to connect two IBM Cloud VPCs so that they can seamlessly communicate with each other, which service can be used?

A. Direct Link

B. Transit Gateway

C. Public Gateway

D. Security Gateway

Answer B. A Transit Gateway can connect two IBM Cloud VPCs so that they can seamlessly communicate with each other.

Question 8.

A company has a new eCommerce application in IBM Cloud VPC. Which service will allow customers to connect directly to a virtual server instance while they are ordering products off the internet?

- A.   Direct Link
- B.   Transit Gateway
- C.   Public Gateway
- D.   Floating IP

Answer D. A Floating IP will allow customers to connect directly to a virtual server instance while they are ordering products off the internet.

Question 9.

In IBM Cloud VPC, _____ is the minimum number of zones in a multizone region.

    A.   2

    B.   3

    C.   4

    D.   5

Answer B. In IBM Cloud VPC, three is the minimum number of zones in a multizone region.

Course 3: Advanced Networking

Introduction and Objectives

In Course 3 of the study guide, the subject matter:

- Introduces advanced networking topics, including load balancers and the use of the IBM Cloud Internet Services (CIS).
- Demonstrates how the networking concepts in this curriculum work together to assist in creating a robust, secure cloud network.

Lessons

- Advanced Networking
- Putting It All Together
- Summary
- Knowledge Check

Objective

- Identify IBM Cloud network security services and options

Course 3: Advanced Networking

Load Balancers

Load balancers distribute requests among servers, reducing traffic flow interruptions. This improves uptime and response time because it reduces the risk of a server being overloaded, and they prevent traffic being sent to a dead server.

Load balancers can be either hardware- or software-based:
- **Hardware load balancers** are physical hardware that require proper management and maintenance. They are less flexible and scalable than software load balancers.
- **Software load balancers** are more flexible, cost-effective, and easier to deploy.

IBM offers several load balancing solutions for clients, including:

IBM Cloud Load Balancer	Application Load Balancer for VPC	Network Load Balancer for VPC
Traffic is distributed among multiple application server instances. Traffic is only forwarded to healthy instances. This ensures that business critical applications are highly available. The IBM Cloud load balancer can be used with Classic infrastructure, including virtual servers and bare metal servers.	For an application load balancer (ALB) for virtual private cloud (VPC), incoming and outgoing packets flow through the load balancer. It provides layer 7 and layer 4 load balancing. A layer 7 load balancer, unlike a layer 4 load balancer, disconnects network traffic and processes the message inside. Secure Sockets Layer (SSL) offloading is also supported.	Public network traffic is sent to the network load balancer (NLB) for VPC. Traffic is then forwarded to target virtual machines (VMs). Only layer 4 load balancing is provided. SSL offloading is not supported.

Course 3: Advanced Networking

IBM Cloud Internet Services

IBM Cloud Internet Services (CIS) is a set of powerful services available to any IBM client using IBM Cloud. CIS assists clients in creating secure, reliable, and high-performing networks.

CIS is powered by Cloudflare. Cloudflare's focus is to ensure everything connected to the internet is secure, private, fast, and reliable. They provide content delivery network (CDN) services, cloud cybersecurity, and distributed denial-of-service (DDoS) mitigation.

The following are options provided by CIS to create a robust network.

Security

Among the many security features CIS offers, the assistance provided in the management of the following may be of particular importance to an organization.

- **Transport Layer Security** CIS can assist clients in managing their transport layer security (TLS). TLS is a cryptographic protocol that protects data being sent between applications over the internet.

- **Web Application Firewalls (WAFs)** CIS assists clients in setting up a WAF. A WAF protects web applications by filtering and monitoring HTTP traffic between a web application and the internet. WAFs operate through a set of rules called policies. WAF policies can be quickly modified and implemented using CIS. Quickly altering policies allows a WAF to respond quickly to DDoS attacks.

- **IP Firewalls** CIS offers several tools for controlling traffic to protect domains, URLs, and directories against volumes of traffic, certain groups of requesters, and particular requesting IPs, such as IPs with a particular address, IP range, or from a specific country.

Reliability

CIS makes use of a global load balancing service to distribute client traffic across multiple servers in multiple regions.

Performance

CIS provides the functionality to assist clients in setting up caching and routing rules for networks and applications.

Course 3: Advanced Networking

Connecting to IBM Cloud Services

When using IBM Cloud services such as Cloud Object Storage, artificial intelligence (AI), or databases, it is often necessary to connect those services to compute resources, such as virtual servers or Red Hat® OpenShift®. By default, the connections go over the IBM Cloud public network. But in some instances, clients may prefer to connect directly to those resources over the private IBM Cloud backbone, without using the public network.

There are two options for doing so:

Cloud Service Endpoints
To connect properly, virtual routing and forwarding (VRF) must be enabled. This moves the IP routing for the account and all of the associated resources to a separate routing table. Once VRF is enabled, IBM Cloud service endpoints may be enabled. When IBM Cloud service endpoints are enabled in an account, clients can choose to expose a private network endpoint when they create a resource. They can then connect directly to this endpoint over the IBM Cloud private network instead of connecting over the public network. Resources that use private network endpoints don't have an internet-routable IP address, so connections to these resources are more secure.

Virtual Private Endpoints
Virtual Private Endpoints (VPEs) are an evolution of Cloud Service Endpoints. VPEs can be used when using resources in a virtual private cloud (VPC). They are virtual IP interfaces that are bound to an endpoint gateway created on a per service, or service instance, basis. The endpoint gateway is a virtualized function that scales horizontally, is redundant and highly available, and spans all availability zones of the client's VPC. Endpoint gateways enable communications from virtual server instances within a client's VPC and IBM Cloud services on the private backbone. VPEs for VPC gives clients the experience of controlling all the private addressing within their cloud.

Course 3: Advanced Networking

IBM Cloud Classic versus IBM Cloud VPC Network

Both IBM Cloud Classic and IBM Cloud VPC network infrastructures can be used to facilitate a robust IBM Cloud network, but the components and mechanics of how they do this are very different. The table here compares the different components and how they function as part of each network type.

Category	Classic	VPC
Location	Data center, pods, virtual local area networks (VLANs)	Regions
Network Security	Physical and virtual appliances from multiple vendors	Cloud-native network functions such as security groups and access control lists (ACLs)
Gateway Routing	Use a virtual or physical network appliance (Vyatta, Juniper vSRX, Fortinet FSA)	Use public gateways and floating IPs
Network Address Translation (NAT)	Use a virtual or physical network appliance (Vyatta, Juniper vSRX, Fortinet FSA)	Bring your own IP (BYOIP) functionality
Load Balancing	IBM Cloud Load Balancer	Appliance Load Balancer Network Load Balancer
Global Load Balancer	IBM Cloud Internet Services	IBM Cloud Internet Services

Question 1.

_____ load balancers are less flexible and scalable.

A.   Hardware

B.   Software

C.   Regional

D.   Global

Answer A. Hardware load balancers are less flexible and scalable.

Question 2.

The _____ feature of IBM Cloud Internet Services (CIS) helps protect web applications by filtering and monitoring HTTP traffic between the application and the internet.

A.   Transport Layer Security (TLS)

B.   Web Application Firewalls (WAFs)

C.   IP Firewalls

D.   SSL Certificates

Answer B. The Web Application Firewalls (WAFs) feature of IBM Cloud Internet Services (CIS) helps protect web applications by filtering and monitoring HTTP traffic between the application and the internet.

Question 3.

Which IBM Cloud option requires virtual routing and forwarding (VRF) be enabled to be able to connect to a service properly?

A. IP routing tables

B. Public network connections

C. Service endpoints

D. Private network endpoints

Answer C. Service endpoints require virtual routing and forwarding (VRF) be enabled to be able to connect to a service properly.

Question 4.

What would be the reason to use Virtual Private Endpoints (VPEs) in a Virtual Private Cloud (VPC)?

A.   Provides virtual IP interfaces for public internet access

B.   Enables private communication between virtual server instances and IBM Cloud services

C.   Controls the public addressing within a client's cloud environment

D.   Acts as redundant and highly available gateways for Virtual Private Cloud resources

Answer B. A reason to use VPEs in a VPC is to enable private communication between virtual server instances and IBM Cloud services.

Question 5.

Network Address Translation in IBM Cloud Classic could use which component(s)?

A. Red Hat OpenShift Network

B. Bring Your Own IP (BYOIP) functionality

C. Virtual or physical network appliance (Vyatta, Juniper vSRX, Fortinet FSA)

D. IBM Cloud service endpoints

Answer C. The components Network Address Translation in IBM Cloud Classic could use are virtual or physical network appliance (Vyatta, Juniper vSRX, Fortinet FSA).

Question 6.

_____ infrastructure uses a virtual or physical network appliance for gateway routing.

A. IBM Cloud VPC Network

B. IBM Bare Metal

C. Red Hat OpenShift Network

D. IBM Cloud Classic

Answer D. The IBM Cloud Classic infrastructure uses a virtual or physical network appliance for gateway routing.

Question 7.

IBM Cloud Internet Services provides what benefits?

A. Improved network performance, security, and reliability

B. Cost-effective cloud storage solutions

C. Enhanced data analytics capabilities

D. Advanced machine learning algorithms for predictive maintenance

Answer A. Improved network performance, security, and reliability are the benefits of using IBM Cloud Internet Services.

Acronyms

Acronym	Acronym Expansion
ACL	Access Control List
ALB	Application Load Balancer
BYOIP	Bring Your Own IP
CIDR	Classless Inter-Domain Routing
CIS	Cloud Internet Services
CDN	Content Delivery Network
DaaS	Desktop-as-a-Service
DDoS	Distributed Denial-of-Service
DSR	Direct Server Return
FSA	FortiGate Security Application
GPO	Group Policy Object
GRE	Generic Routing Encapsulation
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
IAC	Infrastructure as Code
IOPS	Input/Output Operations Per Second
IOS	Organization for Standards
IP	Internet Protocol
LAN	Local Area Network
NAT	Network Address Translation
NLB	Network Load Balancer

Acronyms

Acronym	Acronym Expansion
NTP	Network Time Protocol
RPO	Recovery Point Object
RTO	Recovery Time Object
SFTP	Secure File Transfer Protocol
SSH	Secure Shell
SSL	Secure Sockets Layer
TLS	Transport Layer Security
VCS	vCenter Server
VLAN	Virtual Local Area Network
VPC	Virtual Private Cloud
VPE	Virtual Private Endpoint
VPN	Virtual Private Network
VRF	Virtual Routing and Forwarding
VSI	Virtual Server Instance
WAF	Web Application Firewall
WAN	Wide Area Network
WLAN	Wireless Local Area Network

Courses and Objectives

Courses

1. IBM Cloud Compute Options
2. IBM Cloud Compute Option Use Cases
3. Containers Overview
4. IBM Cloud Container Services

Objectives

- Identify compute options on IBM Cloud® and their advantages
- Identify use cases of compute offerings
- Identify containers, their uses, and their advantages
- Identify container orchestration options on IBM Cloud and their advantages

Course 1: IBM Cloud Compute Options

Introduction and Objectives

In Course 1 of the study guide, the subject matter:

- Provides an overview of the IBM Cloud Bare Metal offerings, including IBM Cloud for VMware® solutions.

- Provides an overview of virtual servers, virtual private cloud (VPC), Power Systems Virtual Servers, Code Engine, confidential computing, and Wazi as a Service offerings.

Lessons

- IBM Cloud Bare Metal Servers
- IBM Cloud for VMware Solutions
- IBM Cloud Virtual Servers
- IBM Cloud Power Systems Virtual Servers
- Confidential Computing
- IBM Wazi as a Service
- Summary
- Knowledge Check

Objectives

- Identify compute options on IBM Cloud and their advantages
- Differentiate between IBM Cloud Virtual Server Instances (VSIs) and bare metal server offerings

Course 1: IBM Cloud Compute Options

IBM Cloud Bare Metal Servers

IBM Cloud Bare Metal servers are physical, single-tenant, dedicated servers that provide ultimate performance and control with low-level access to the hardware resources. Bare Metal servers can be included on either the IBM Cloud Classic or IBM Cloud Virtual Private Cloud (VPC) network.

Benefits of Bare Metal Servers
- Physical isolation and therefore a great quality of service (QoS)
- Powerful processing
- Consistent disk and network input/output (I/O) performance
- Seamlessly scaled workloads
- Discount over on-demand pricing

IBM Cloud Bare Metal Servers on Classic offers clients the ability to choose from tens of thousands of configurations. The flexibility in configuration allows clients to choose the best compute match for the planned workload.

IBM Cloud Bare Metal Servers on VPC options allow clients to deploy servers in 10 minutes or less. This option allows clients to enjoy the security and performance of the private cloud with the flexibility and scalability of the public cloud. This is a good option for clients focused on quick provisioning times and fast networking speeds.

Course 1: IBM Cloud Compute Options

VMware Offerings

Make it simpler for organizations with an investment in VMware to make the most of the potential of cloud computing.

IBM Cloud VMware solutions assist both clients who want complete control and management of their VMware environment and clients who do not want to manage the underlying infrastructure and VMware hypervisor.

VMware Solutions:
- **VMware as a Service**
  - Provides the VMware Cloud Director platforms as a managed service
  - IBM performs the configuration, hosting, operations, and life cycle management of the VMware software, so clients can quickly deploy their VMware-based cloud computing environments.
- **VMware Solutions Shared**
  - Provides standardized and customizable deployment of VMware virtual data center environments
  - Provides a self-service on-demand VMware cloud computing platform with VMware Cloud Director running on IBM Cloud
  - Provides deployment options:
    - Multitenant on-demand virtual data centers
    - Multitenant reserved virtual data centers
- **VMware vSphere**
  - Assists clients who would like to manage and control their VMware environment
  - Provides a streamlined and optimized ordering platform for VMware
  - Allows clients to build their own IBM-hosted VMware environment by customizing and ordering the VMware-compatible hardware
  - Does not automate the installation, configuration, or bring up optional VMware components

Course 1: IBM Cloud Compute Options

VMware Offerings, Continued

VMware Solutions, Continued:
- **VMware vCenter Server**
  - Is a hosted private cloud that delivers the VMware vSphere stack as a service. The VMware environment is built, in addition to a minimum of three IBM Cloud Bare Metal Servers
  - Can rapidly and elastically scale the compute capacity up and down as needed
- **VMware Regulated Workloads**
  - Offers a secure-by-default architecture that follows IBM's unique policy controls framework
  - Offers the highest level of data encryption and continuous compliance monitoring
  - Is available for both single zone and multizone regions
  - Provides key management services such as Bring Your Own Key (BYOK) and Keep Your Own Key (KYOK)
- **IBM Cloud Cyber Recovery**
  - Is an option for organizations looking for complete cyber resilience solutions to help protect from cyberthreats, cyberattacks, and ransomware

The chart displayed here provides a quick comparison and contrast for all the VMware platform services covered for both Classic and VPC networking environments.

	VMware vSphere self-managed service platform			VMware Cloud Director IBM-managed service platform	
Solution	VMware vCenter Server (VCS) on Classic	VMware vSphere on Bare Metal Servers for Classic	VMware vSphere Bare Metal Servers for VPC	VMware as a Service (VMWaaS) on Classic	VMware Shared on Classic
Summary	IBM automation deploys and configures; client maintains/monitors VMware	Client designs, deploys, configures, and maintains/monitors VMware	Client designs, deploys, configures, and maintains/monitors VMware	IBM automation deploys and configures, and IBM also maintains and monitors VMware	IBM automation deploys and configures, and IBM also maintains and monitors VMware and manages compute capacity
Control	Highest with root access to hypervisor	Highest with root access to hypervisor	Highest with root access to hypervisor	No root access to hypervisor	No root access to hypervisor
Optional IBM-provided service	Cyber Recovery	Cyber Recovery	Cyber Recovery	Cyber Recovery	Cyber Recovery

Course 1: IBM Cloud Compute Options

IBM Cloud Virtual Servers

Virtual servers are another compute option for IBM Cloud. As with bare metal servers, virtual servers may run on IBM Cloud Classic infrastructure or on IBM Cloud VPC infrastructure.

Virtual Servers for Classic Infrastructure

Virtual servers in Classic infrastructure may be created in either a public or dedicated environment. Virtual servers are scalable, can be deployed in minutes, and clients have a choice of virtual server images and geographic regions.

Features of virtual servers in Classic infrastructure include seamlessly integrated, fully customized, rapidly provisioned, remotely managed, and easily scalable.

Deployment Options

Virtual servers may be deployed on:

- **Public Virtual Servers** are designed to help clients quickly get up and running. They are billed monthly or hourly and are deployed on multitenant hosts.
- **Transient Virtual Servers** are best for flexible workloads and for clients that want to reduce costs. They take advantage of unused capacity; however, transient virtual servers can be de-provisioned at any time by IBM Cloud without notification. Also, transient virtual servers cannot be upgraded, downgraded, or use local storage during the contracted term.
- **Reserved Virtual Servers** allow the client to reserve up to 20 virtual server instances and provision them whenever necessary (under a set contract term).
- **Dedicated Virtual Servers** have two deployment models: dedicated hosts (offers help with control over workload placement) and dedicated instances (offers single-tenant isolation). By using dedicated virtual servers, clients can specify the affinity and anti-affinity rules, whether it's host to virtual server or virtual server to virtual server.

Course 1: IBM Cloud Compute Options

IBM Cloud Virtual Servers, Continued

Comparing Bare Metal Servers and Virtual Servers

The table displays the advantages and strengths of bare metal servers and virtual servers.

Bare Metal Servers	Virtual Servers
Ideal for CPU-and-Input/Output heavy workloads	Quickly provisioned
Offers the best performance and security	Flexible and scalable environment
Flexible, transparent, and offers control	Cost-effective
Cannot suspend billing	Provides the ability to suspend billing

Course 1: IBM Cloud Compute Options

Power Systems Virtual Servers

- Are a Power Systems offering
- Are also known as a logical partition (LPAR), run on IBM Power Systems hardware with PowerVM® hypervisor
- Are located next to IBM Cloud data centers
- Have separate networks and direct-attached storage, with the ability to connect back to IBM Cloud data centers
- Can run AIX®, IBM i, or Linux® operating systems
- Benefit: Makes it simple for clients with Power Systems workloads to transition services to the cloud or to quickly build development and test environments
- Key Features: Customized infrastructure, two tiers of storage available, images from AIX and IBM i, and monthly billing rates

Power Systems Virtual Servers

- Applications
- Data
- Runtime
- Middleware
- Operating system
- Virtualization
- Servers
- Storage
- Networking

This graphic represents a responsibility assignment (RACI) matrix for Power Systems Virtual Servers. Client responsibilities are in grey. IBM Cloud responsibilities are in blue.

Course 1: IBM Cloud Compute Options

Confidential Computing

Confidential computing technology fills a key gap in overall data security. Encryption services help to protect data at rest (in other words, data in storage and in databases) and in transit (in other words, data moving over a network), but data in memory must be unencrypted prior to being processed. When data is unencrypted in memory, it is exposed to memory dumps and other malicious attacks.

Confidential computing fills this gap by isolating data in a protected central processing unit (CPU) as it is processed. Both the data being processed and the techniques used to process it are accessible only to authorized programming code.

Confidential computing leverages a hardware-based trusted execution environment (TEE). The TEE is secured using embedded encryption keys, which are only accessible to authorized application code. If unauthorized code attempts to hack or alter the keys, the TEE denies access to the keys and cancels the computation.

Confidential computing environments:

- IBM z15 and IBM LinuxONE III hardware provide the architecture for confidential computing.
- IBM Secure Execution for Linux technology enables creation of a TEE.
- IBM Secure Execution for Linux technology is used to create an IBM Cloud Hyper Protect Virtual Server for VPC instance.

IBM Wazi as a Service

The IBM Wazi as a Service solution assists organizations in managing and modernizing applications on IBM Z® Systems platforms using IBM Cloud. It enables development and testing of z/OS® application components in a virtual infrastructure running on IBM Z architecture on IBM public cloud. IBM Wazi as a Service provides cloud security, the ability to dynamically scale z/OS resources where and when they are needed, and integration of z/OS resources with public cloud services. Three basic components of IBM Wazi as a Service:
- **z/OS Dev and Test Systems** allow developers to quickly spin up a z/OS development and test system in IBM Cloud VPC.
- **Wazi Image Builder** provides enterprises the ability to create custom images from an on-premises logical partition (LPAR) and deploy these images as custom images within the IBM Cloud VPC infrastructure.
- **z/OS templates** extend the IBM Continuous Delivery Service to automate builds in secure DevSecOps toolchains.

Question 1.

The VMware option, _____ , on IBM Cloud includes automated deployment and configuration in a dedicated environment.

A. VMware vCenter Server

B. VMware vSphere

C. VMware Shared

D. VMware ESXi

Answer A. The VMware option, VMware vCenter Server, on IBM Cloud includes automated deployment and configuration in a dedicated environment.

Question 2.

Which IBM Cloud VMware Solution should an organization use that encrypts both data and virtual machines in addition to **providing** a key management service?

A.   VMware vSphere

B.   VMware Regulated Workloads

C.   VMware Cyber Recovery

D.   VMware vCenter Server

Answer B. The VMware Regulated Workloads solution is what an organization should use that will encrypt both data and virtual machines in addition to providing a key management service.

Question 3.

Which IBM Cloud for VMware solution could a client use who wants to quickly deploy their VMware-based cloud computing environments without concerns about configuration, hosting, operations, and life cycle management of the VMware software?

A.   VMware vCenter Server

B.   VMware as a Service

C.   VMware vSphere

D.   VMware Regulated Workload

Answer B. VMware as a Service could be used by a client who wants to quickly deploy their VMware-based cloud computing environments without concerns about configuration, hosting, operations, and life cycle management of the VMware software.

Question 4.

The _____ IBM Cloud for VMware solution offers a multitenant deployment option.

A. VMware vCenter Server

B. VMware Shared

C. VMware vSphere

D. VMware Regulated Workload

Answer B. The VMware Shared IBM Cloud for VMware solution offers a multitenant deployment option.

Question 5.

What option would allow a startup with a limited budget to sacrifice up time for cost savings?

    A.   Public Virtual Servers

    B.   Dedicated Virtual Servers

    C.   Reserved Virtual Servers

    D.   Transient Virtual Servers

Answer D. Transient Virtual Servers would allow a startup with a limited budget to sacrifice up time for cost savings.

Question 6.

IBM Cloud _____ allow specifying the affinity and anti-affinity rules related to a physical machine or rack.

A.   Public Virtual Servers

B.   Dedicated Virtual Servers

C.   Reserved Virtual Servers

D.   Transient Virtual Servers

Answer B. IBM Cloud Dedicated Virtual Servers allow specifying the affinity and anti-affinity rules related to a physical machine or rack.

Question 7.

When comparing bare metal servers with virtual servers on IBM Cloud, what is a benefit of bare metal servers?

A. Scalable environment

B. Ideal for CPU and I/O heavy workloads

C. Provides the ability to suspend billing

D. Quickly provisioned

Answer B. When comparing bare metal servers with virtual servers on IBM Cloud, a benefit of bare metal servers is that it's ideal for CPU and I/O heavy workloads.

Question 8.

If a client wants to migrate an on-premises LPAR running AIX to IBM Cloud, which type of service should they migrate to?

A. IBM VMware Shared Solutions

B. VSIs on IBM Classic

C. IBM Power Virtual Servers

D. Bare Metal Servers for VPC

Answer C. A client who wants to migrate an on-premises LPAR running AIX to IBM Cloud should migrate to IBM Power Virtual Servers.

Question 9.

What is IBM responsible for based on the shared responsibility model?

A.   Virtualization
B.   Operating System
C.   Middleware
D.   Runtime

Answer A. IBM is responsible for virtualization in the shared responsibility model when working with Power System Virtual Servers.

Question 10.

The _____ technology is what enables confidential computing for IBM Cloud Hyper Protect virtual servers.

A. XEN

B. KVM

C. TXT

D. TEE

Answer D. The TEE technology is what enables confidential computing for IBM Cloud Hyper Protect virtual servers.

Question 11.

The _____ solution enables development and testing of z/OS application components in a virtual infrastructure running on IBM Z architecture on IBM Cloud.

A.   Red Hat OpenShift on IBM Cloud

B.   VMware Solutions

C.   IBM Wazi as a Service

D.   Z Systems Virtual Servers

Answer C. The IBM Wazi as a Service solution enables development and testing of z/OS application components in a virtual infrastructure running on IBM Z architecture on IBM Cloud.

Course 2: IBM Cloud Compute Use Cases

Introduction and Objectives

In Course 2 of the study guide, the subject matter:

- Reviews the use cases of the various IBM Cloud compute options including bare metal servers, VMware, and virtual servers.

- Explores the nuances of the Skytap cloud service, and finally, delves into virtual desktop applications, confidential computing, and Wazi as a Service.

Lessons

- Bare Metal Use Cases on IBM Cloud
- VMware Use Cases on IBM Cloud
- Virtual Server Use Cases on IBM Cloud
- Skytap Service on IBM Cloud
- IBM Cloud Virtual Desktop Applications
- Confidential Computing Use Cases on IBM Cloud
- IBM Wazi as a Service Use Cases
- Summary
- Knowledge Check

Objective

- Identify the use cases of compute offerings

Course 2: IBM Cloud Compute Use Cases

Bare Metal Use Cases
- Typically used where performance is a key concern, such as where data needs to be retrieved quickly from local storage or where a specific hardware type is required for heavy workloads
- Helpful when a dedicated environment is needed due to security, policy, or compliance requirements
- Useful for games and other applications that use heavy bandwidth and require high-level performance

VMware Use Cases
Allows companies to extend and migrate VMware workloads to the cloud, including:
- Data center transformation
- Disaster recover
- Security and compliance
- Reuse existing skill sets when going from on premises to the cloud

Virtual Server Use Cases
Allows companies to reduce server hardware expenses as well as physical power and energy costs by sharing resources through running multiple, yet distinct, servers on a single physical machine.
Use Cases for Classic and VPC:
- Rapid provisioning: Virtual servers are provisioned in as few as five minutes.
- Easy Scalability: After provisioning, a virtual server can be quickly and easily scaled up or down.
- Billing Suspension: Billing stops automatically when the instance is stopped.
Other Use Cases:
- High network bandwidth
- High availability built in

Skytap Service on IBM Cloud Use Cases
- Can easily migrate AIX, IBM i, Linux, and Windows and can be quickly provisioned
- Highly available and can create a self-managed application environment

Course 2: IBM Cloud Compute Use Cases

Virtual Desktop Application for IBM Cloud Offering

Desktop applications for the cloud environment deliver complete virtual desktop environments to users from the cloud. These environments include operating systems, applications, files, and user preferences.

Some virtual desktop applications available as services in IBM Cloud include:

- **Citrix DaaS for IBM Cloud**
    - Is primarily a client-managed solution
- **Dizzion DaaS**
    - Is a managed virtual desktop service that runs VMware Horizon Enterprise
    - Dizzion maintains the underlying VMware infrastructure, and the client provides a desktop image for Dizzion to deploy.

Hyper Protect Virtual Servers and Confidential Computing on IBM Cloud Use Cases

IBM protects data in use through Hyper Protect Virtual Servers. These servers make it possible to bring IBM Z into IBM's global public cloud data centers.

Examples of when Hyper Protect Virtual Servers can be used to bring the power of confidential computing to an organization:

- Run both core and non-core workloads in a public cloud, observing all security and compliance policies of an enterprise.
- Protect the assets of a business while simultaneously maintaining enhanced business service levels.
- Instantiate Linux virtual servers with the organization's own public SSH key.
- Deploy any supported workload into the most secure, highly performant Linux system.

IBM Wazi as a Service Use Cases

IBM Cloud Wazi as a Service accelerates cloud native development and testing of z/OS applications with z/OS Virtual Server in IBM Cloud VPC. It's a protected space in IBM Cloud with the security of a private cloud and the agility of a public cloud.

Examples of when clients might wish to engage Wazi as a Service:

- Development and testing
- Infrastructure testing
- Innovation and skill building
- Secure z/OS Apps with DevSecOps practices.

Question 1.

A software company is interested in migrating their workloads to IBM Cloud and have the following compute requirements: high performance, dedicated compute, and high customization. Which compute option meets all of these requirements?

A.   Bare Metal Server

B.   VMware Solutions Shared

C.   VSI for VPC

D.   Code Engine

Answer A. The bare metal server option will meet these requirements.

Question 2.

A gaming startup company requires dedicated and secure compute capacity to deploy and test a new game they have developed. The game is network intensive and demands high performance from the compute. Based on company expectations, what would be the most appropriate compute option to use?

A.   Virtual servers with high compute capacity

B.   Virtual servers with high network bandwidth

C.   Bare metal servers

D.   Virtual servers with high compute and network bandwidth

Answer C. Bare metal servers would be the most appropriate compute option to use.

Question 3.

A client has a large investment in VMware software and skills. They are looking for a route to easily migrate and manage services in the cloud. Which IBM Cloud compute option would help with the least disruption?

A. Red Hat OpenShift

B. VSIs on VPC

C. VMware Solutions

D. IBM Cloud Kubernetes Service

Answer C. VMware Solutions would be the best option for this client.

Question 4.

A client is looking for a solution in the cloud. They have two requirements: high computing capacity during specific daily intervals and to only pay for the time they are using the service. What is the best IBM Cloud offering to recommend to this client?

A.   Bare Metal for VPC

B.   VMware Solutions

C.   VMware on VPC

D.   VSIs for VPC

Answer D. VSIs for VPC is the best IBM Cloud offering to recommend to this client.

Question 5.

What would be a benefit for a client to choose IBM Wazi as a Service on IBM Cloud?

A.   Virtual storage in cloud

B.   Use of virtual desktops

C.   Critical production workload

D.   Development and testing

Answer D. Development and testing would be a benefit to a client if they used IBM Wazi as a Service on IBM Cloud.

Question 6.

What is a reason a company would use IBM Cloud VPC z/OS virtual machines?

A. To start a free development and test environment in any public cloud

B. To start a development and test environment in minutes without dependency in operations

C. To start a development and test environment in a local MZR

D. To start a development and test environment on premises

Answer B. To start a development and test environment in minutes without dependency in operations is a reason a company would use IBM Cloud VPC z/OS virtual machines.

Question 7.

Operating systems and applications are provided in a virtual desktop environment offered on IBM Cloud. What else is offered?

A. Files and user preferences

B. Bare metal servers and cloud access

C. User data configuration and cloud access

D. Storage, memory, and high-speed internet access

Answer A. Operating systems, applications, files, and user preferences are provided in a virtual desktop environment offered on IBM Cloud.

Course 3: Containers Overview

Introduction and Objectives

In Course 3 of the study guide, the subject matter:

- Introduces container orchestration systems.

- Explores containers and how they differ from virtual machines, as well as different methods for handling containers.

Lessons

- Containers and Virtual Machines

- Containers, Container Orchestration, and Kubernetes

- IBM Cloud Code Engine

- Horizontal and Vertical Scaling for Containers

- Summary

- Knowledge Check

Objective

- Identify containers, their uses, and their advantages

Course 3: Containers Overview

Containers

Containers are self-contained units that are designed to hold applications as well as all libraries and dependencies that the application requires. Containers can be run on physical computers or in the cloud.

Containers do not have an operating system of their own. Instead, they access the operating system of the device they are running on using a hypervisor. This allows them to be more lightweight, easier to provision, and much smaller than traditional virtual machines. It also allows them to remain isolated from each other.

Containers versus Virtual Machines (VMs)

VMs

Traditional virtualization leverages a hypervisor to virtualize physical hardware.
Each VM contains:
- A guest operating system (OS)
- A virtual copy of the hardware that the OS requires to run
- An application and its associated libraries and dependencies

Containers

Containers, on the other hand, virtualize the operating system. Each container only contains:
- The application
- Its libraries
- Its dependencies

Course 3: Containers Overview

Containers and Container Orchestration Systems

Containerization
- The process of designing and packaging software in containers is called containerization.
- When an application is containerized, it is packaged with its relevant environment variables, configuration files, libraries, and software dependencies.
- This container image can then be run on a container platform.

Container Orchestration
- Container orchestration is a way to manage large volumes of containers.

Orchestration Platform
- Kubernetes is a widely used, open source software that is designed specifically as a container orchestration platform.
- It does not limit the types of applications that can run on it.

Use Cases for Containers
- **Microservices**: Applications are constructed of many loosely coupled and independently deployable smaller services.
- **DevOps**: The combination of a microservices architecture and containers is a common foundation for many teams that embrace DevOps as the way they build, ship, and run software.
- **Hybrid, multivendor**: Since containers can run across laptops, on premises, and cloud environments, this makes them an ideal underlying architecture for hybrid cloud and multi-cloud (or multi-vendor) scenarios, where organizations operate across multiple public clouds in addition to their own data centers.
- **App Modernizing and migration**: App modernization and migration to the cloud is one example of using containers to replace virtual machines.

Course 3: Containers Overview

Code Engine

Code Engine is a fully-managed serverless platform. It runs container-based workloads and provides a friendly development platform. It scales in response to load without the overhead of configuring servers. Code Engine provides the ability to run containers for any type of workload, including serverless, batch jobs, and even traditional web applications.

Benefits of Code Engine

Benefits of Code Engine include:
*   Friendly development platform
*   Ability to scale easily without having to worry about configuring servers
*   Fast deployment
*   Fully integrated into IBM Cloud
*   IBM-managed infrastructure
*   Code Engine supports both Dockerfile and Cloud Native Buildpack. It pulls the source code and creates the container image.
*   Provides privacy
*   Allows for event-based reactions
*   Auto scales workloads

Security

Code Engine is built with security in mind. Clients and their corresponding workloads are isolated from one another. This is due to projects which are based on Kubernetes namespaces. Role-based access controls occur on a resource level, which allows only authorized users to perform certain operations.

Horizontal and Vertical Scaling for Containers

When a client is running applications or workloads on a VM, or a local partition, either on-premises or on the cloud, they may run out of memory (OOM). They can scale up or scale out.

Scaling Up: Clients can quickly scale up simply by increasing the Central Processing Unit (CPU) and Random Access Memory (RAM). One benefit of this option is that they do not need to purchase any additional servers. However, when scaling up, it can mean that they will have some downtime as they resize their VM.
Scaling Out: Clients can scale out by provisioning additional VMs to spread out the workload. While this does create a more complicated environment, and load balancers and databases will need to be configured, it enables them to scale their environments for improved performance.

Question 1.

What would be considered the main difference between containers and virtual machines?

A.   Containers virtualize physical hardware, while virtual machines virtualize the operating system.

B.   Containers include only the application, libraries, and dependencies, while virtual machines contain a guest operating system, virtual hardware, and application.

C.   Containers have a larger footprint than virtual machines, making them faster to recreate and restart.

D.   Containers and virtual machines are identical in terms of their structure and functionality.

Answer B. Containers include only the application, libraries, and dependencies, while virtual machines contain a guest operating system, virtual hardware, and application.

Question 2.

When operating large volumes of containers, what would be an ideal characteristic?

A. Compatibility with legacy applications

B. High security features

C. Lightweight and small size

D. Ability to store large amounts of data

Answer C. When operating large volumes of containers, lightweight and small size would be an ideal charateristic.

Question 3.

The _____ architectural style is the best option to use when constructing applications consisting of loosely coupled and independently deployable smaller services.

A. Service-oriented

B. Monolithic

C. Microservices

D. Client-server

Answer C. The Microservices architectural style is the best option to use when constructing applications of loosely coupled and independently deployable smaller services.

Question 4.

When using IBM Cloud Code Engine to run a cloud native application, IBM Cloud Code Engine could be described as _____.

A. A development platform for traditional web applications

B. A platform for configuring and managing servers for container-based workloads

C. A serverless platform for running container-based workloads

D. A batch processing platform for serverless workloads

Answer C. An IBM Cloud Code Engine running a cloud native application could be described as a serverless platform for running container-based workloads.

Question 5.

What isolates the workloads between clients in IBM Cloud Code Engine?

A. Resource quotas

B. Role-based access controls

C. Kubernetes namespaces

D. Security groups

Answer C. Kubernetes namespaces isolate the workloads between clients in IBM Cloud Code Engine.

Course 4: IBM Cloud Container Services

Introduction and Objectives

In Course 4 of the study guide, the subject matter:

▪ Introduces two container orchestration software options: IBM Cloud Kubernetes Service (IKS) and Red Hat® OpenShift® on IBM Cloud.

▪ Includes a brief introduction to Red Hat OpenShift and Kubernetes application exposure services.

Lessons

▪ IBM Cloud Kubernetes Service

▪ Red Hat OpenShift on IBM Cloud

▪ Red Hat OpenShift and Kubernetes Application Exposure Services

▪ Summary

▪ Knowledge Check

Objective

▪ Identify container orchestration options on IBM Cloud and their advantages

Course 4: IBM Cloud Container Services

IBM Cloud Kubernetes Service

Kubernetes is an open source platform for managing containerized workloads and services across multiple hosts. It includes tools for deploying, automating, monitoring, and scaling containerized apps with minimal manual intervention. **IBM Cloud Kubernetes Service (IKS)** is a managed offering clients use to create their own Kubernetes clusters to deploy and manage containerized apps on IBM Cloud.

The IKS managed offering can assist clients in using the power of Kubernetes to its fullest in delivering highly available and secure containerized apps to the cloud. When clients engage IKS, IBM takes over the management of the master node, freeing the client from having to worry about the host operating system, container runtime, or Kubernetes versioning and updates.

Kubernetes Basics

A Kubernetes cluster consists of a master and one or more compute hosts, commonly called worker nodes. Worker nodes are organized into worker pools. All workers in the pool have the same profile (CPU, memory, operating system).

Cluster resources are divided into separate areas using **namespaces**. Organizations can deploy apps into a particular namespace.
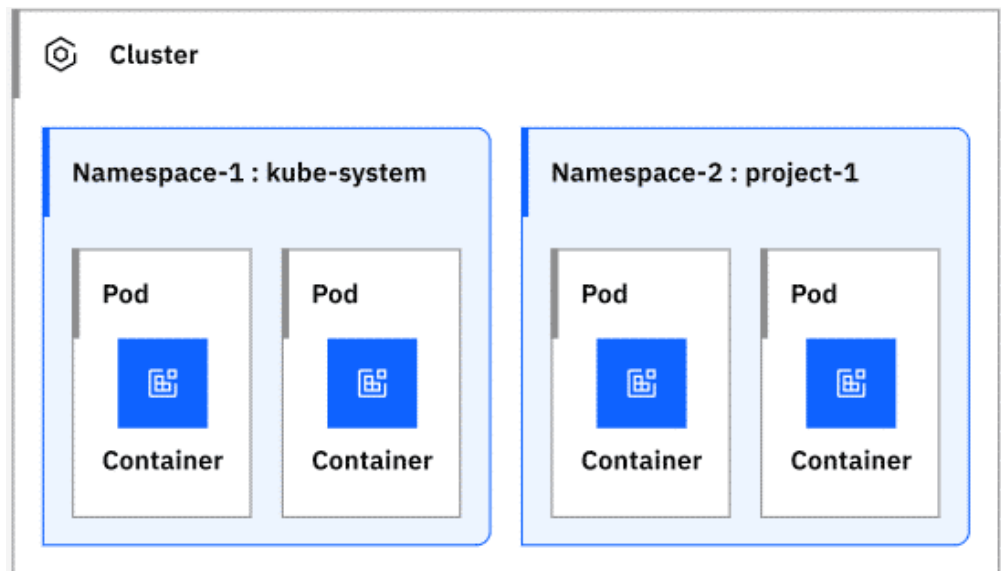
Pods group containers that must be treated as a single unit. Usually, each container is deployed in its own pod. Services provide network connectivity to the pods.

A **container** is a packaged app with all its dependencies (code, run time, system tools, libraries, and settings).

Course 4: IBM Cloud Container Services

Information Storage

IBM stores two types of information for each cluster created with IBM Cloud Kubernetes Service:
- **Personal information**: The email address of the IBM Cloud account that created the cluster.
- **Sensitive information**:
  - The TLS certificate and secret used for the assigned Ingress subdomain
  - The certificate authority used for the TLS certificate
  - The certificate authority, private keys, and TLS certificates for the Kubernetes master components, including the Kubernetes API server, etcd data store, and VPN
  - A client root key in IBM Key Protect for each IBM Cloud account that is used to encrypt personal and sensitive information

All information is stored in an etcd database and backed up every eight hours to IBM Cloud Object Storage. The etcd database and Object Storage service instance are owned and managed by the IBM Cloud SRE team. For each IBM Cloud account, a client root key in IBM Key Protect is created and managed by the IBM Cloud Kubernetes Service team. This root key is used to encrypt all personal and sensitive information in etcd and in Object Storage.

Course 4: IBM Cloud Container Services

Red Hat OpenShift

Red Hat OpenShift is another open source container management platform. Red Hat OpenShift is built using the same underlying technology that is used in IBM Cloud Kubernetes Service; however, Red Hat OpenShift is built on top of the Kubernetes platform and other open source projects such as Docker and Istio. Kubernetes does not include Red Hat OpenShift services, and it is its own standalone option.

Kubernetes and Red Hat OpenShift on IBM Cloud Comparison

The underlying architecture of IKS and Red Hat OpenShift on IBM Cloud is the same. For example, the master nodes for Red Hat OpenShift on IBM Cloud are provisioned in an IBM Cloud account managed by IBM, and the worker nodes are provisioned in the same way as IKS worker nodes on Virtual Private Cloud (VPC) or Classic infrastructure.

The two technologies have some differences. The chart below categorizes the differences into four main categories: deployment, management, operations, and security.

Technology	Deployment	Management	Operations	Security
Kubernetes	It takes a flexible approach to deploying applications. It might be considered a more comprehensive solution because of the flexibility it offers.	Has a built-in dashboard to manage applications but may not be powerful enough for all users	Adding additional resources to the cluster can be very time consuming and requires scripts to be developed.	Role-based attribute controls (RBAC) and identity and access management (IAM) can be built out, but it takes time to do so.
Red Hat OpenShift	There is a more guided and structured process available. It might be considered a more attractive option because of the automation it provides.	Has a prescribed way of managing applications	Ansible playbooks and installers are available to bring new resources into the cluster. There are also ways to implement autoscaling.	RBAC and IAM are built-in. It's designed for regulatory-readiness and meets GDPR and HIPAA compliance regulations.

Course 4: IBM Cloud Container Services

Connecting Containers to External Traffic

When using IBM Cloud OpenShift or Kubernetes container platforms, clients may need to allow external traffic to the application running in the service. A few capabilities are built into the service that provide the ability to securely expose these applications to external traffic.

OpenShift Ingress Controller

OpenShift Ingress controller can expose multiple apps in a cluster. The Ingress subdomain is used as a secure and unique public or private entry point to route incoming requests. One subdomain can be used to expose multiple apps in the cluster as services.

Ingress exposes apps to external traffic through the OpenShift Ingress controller. The OpenShift Controller Manager converts the Ingress resources to the Route resources. The OpenShift Ingress controller is responsible for processing the routes. The Ingress Controller can also customize routing rules. It can also implement custom routing rules and SSL termination.

OpenShift Router

The OpenShift router is a HAProxy-based Kubernetes service. It manages incoming traffic for the apps in the cluster. This occurs through implementing routing rules for the apps. The router finds incoming HTTP or HTTPS service requests. These requests are forwarded to the pods for that app only according to the rules defined in the router resources.

NodePort

NodePorts are best for testing public or private access or providing access for a small amount of time.

Load Balancer

Load balancer depends on the cluster's infrastructure provider:
- **Classic clusters** use network load balancer (NLB). The NLB can be customized by exposing any port that the app requires. Public NLBs can have a subdomain for the app.
- **VPC clusters** use load balancer for VPC. The VPC load balancer routes requests for the app. This occurs through the private NodePorts, which are automatically opened on the worker nodes.

Question 1.

A client is looking to fully leverage the power of Kubernetes in delivering highly available and secure containerized applications. The _____ offering would help this client's needs.

A.   IBM Container Service

B.   Kubernetes Cloud Service

C.   IBM Kubernetes Service

D.   Managed Kubernetes cluster service

Answer C. The IBM Kubernetes Service offering would help a client who is looking to fully leverage the power of Kubernetes in delivering highly available and secure containerized applications.

Question 2.

When using IBM Cloud Kubernetes Service, cluster related information is stored in
_____.

A.   Object Storage

B.   etcd database

C.   Db2 database

D.   Cloudant database

Answer B. Cluster related information is stored in the etcd database when
using IBM Cloud Kubernetes Service.

Question 3.

_____ is responsible for backing up the etcd database in the master cluster of a Red Hat OpenShift on the IBM Cloud service.

A.   IBM Cloud

B.   The client

C.   Red Hat

D.   The partner

Answer A. IBM Cloud is responsible for backing up the etcd database in the master cluster iof a Red Hat OpenShift on the IBM Cloud service.

Course 4: IBM Cloud Container Services
Check Your Knowledge

Question 4.

The IBM Cloud service that is designed to meet GDPR and HIPAA regulatory readiness is _____ .

A.   IBM Cloud Foundry

B.   IBM Cloud functions

C.   Red Hat OpenShift on IBM Cloud

D.   IBM Container service

Answer C. The IBM Cloud service that is designed to meet GDPR and HIPAA regulatory readiness is Red Hat OpenShift on IBM Cloud.

Question 5.

What is the functionality of the Ingress Controller in Red Hat OpenShift on IBM Cloud?

A. Exposes apps to internal traffic

B. Exposes apps to other apps inside the cluster

C. Exposes apps to external traffic

D. Exposes databases to apps

Answer C. Exposing apps to external traffic is the functionality of the Ingress Controller in Red Hat OpenShift on IBM Cloud.

Question 6.

In an IBM Cloud Kubernetes Service Cluster, _____ provides the mechanism for isolating a cluster resource in a logical separation.

A.   kube-proxy

B.   Namespaces

C.   Pods

D.   Secrets

Answer B. In an IBM Cloud Kubernetes Service Cluster, Namespaces provide the mechanism for isolating a cluster resource in a logical separation.

Question 7.

IBM Kubernetes Service is one open source container management platform available on IBM Cloud. What is another open source container management platform **available on IBM Cloud?**

A. Elastic Kubernetes Service

B. IBM Container Service

C. Red Hat OpenShift Service

D. Elastic Container Service

Answer C. IBM Kubernetes Service and Red Hat OpenShift Service are two open source container management platforms available on IBM Cloud.

Acronyms

Acronym	Acronym Expansion
API	Application Programming Interface
BYOK	Bring Your Own Key
BYOL	Bring Your Own License
CPU	Central Processing Unit
CVAD	Citrix Virtual Apps and Desktops
DaaS	Desktop-as-a-Service
GPO	Group Policy Object
HPCR	Hyper Protect Container Runtime
HTTP	Hypertext Transfer Protocol
HTTPS	Hypertext Transfer Protocol Secure
IAM	Identity Access Management
IDE	Integrated Development Environment
IKS	IBM Cloud Kubernetes Service
I/O	Input/Output
IOPS	Input/Output Operations Per Second
KYOK	Keep Your Own Key
LPAR	Logical Partitions
NVMe	Nonvolatile Memory Express
QoS	Quality of Service
RACI	Responsible Assignment

Acronyms

Acronym	Acronym Expansion
RAM	Random Access Memory
RBAC	Role-Based Access
RPO	Recovery Point Object
RTO	Recovery Time Object
SFTP	Secure File Transfer Protocol
SSD	Solid-State Drive
SSH	Secure Shell
SSL	Secure Sockets Layer
TEE	Trusted Execution Environment
TLS	Transport Layer Security
VCS	vCenter Server
VDI	Virtual Desktop Infrastructure
VLAN	Virtual Local Area Network
VM	Virtual Machine
VPC	Virtual Private Cloud
VSI	Virtual Server Instance
WaaS	Wazi as a Service

Courses and Objectives

Courses

1. IBM Cloud Storage Offerings
2. IBM Cloud Managed Databases
3. Storage Security Options
4. IBM Hyper Protect Crypto Services
5. IBM Backup, Recovery, and Replication

Objectives

- List the high-level concepts of IBM Cloud® storage offerings and their advantages
- Recognize the value of the different IBM Cloud-managed database options
- Identify storage security options for IBM Cloud
- List the benefits and options of IBM Hyper Protect Crypto Services
- Recognize the benefits and options of IBM Backup, Recovery, and Replication Solutions

Course 1: IBM Cloud Storage Options

Introduction and Objectives

In Course 1 of the study guide, the subject matter:

▪ Explores the high-level concepts of IBM Cloud® storage offerings, advantages, and use cases.

Lessons

▪ Introduction to Storage Options

▪ Block Storage

▪ File Storage

▪ Object Storage

▪ Software-Defined Storage (SDS) Offerings

▪ Summary

▪ Knowledge Check

Objective

• List the high-level concepts of IBM Cloud storage offerings and their advantages

• Identify use cases for IBM Cloud storage offerings

Course 1: IBM Cloud Storage Options

Storage Options Overview

The three most common types of storage are:

Block

Block storage is a technology that is used to store data files on Storage Area Networks (SANs) or cloud-based storage environments by breaking the data into blocks and storing those blocks as separate pieces.

File

File storage is storage that is typically presented from Network Attached Storage (NAS) technology. Files are typically stored in a directory tree structure using folders, in the same way seen on a local hard drive.

Object

Object storage is a type of unstructured storage, which is great for storing large amounts of data. Objects stored can range from small text files to large data dumps or multimedia files. These are stored in repositories called buckets, which can be replicated for high availability.

Course 1: IBM Cloud Storage Options

Block Storage

Block storage stores data by breaking it into evenly sized blocks, which is written to disk. Used for computing situations where users require fast, efficient, and reliable data storage.

Block Storage Features

Can be deployed on a virtual private cloud (VPC) or classic environment

Endurance and custom tiers

Data stored is highly available because each block is stored multiple times across different disks.

Data immediately becomes inaccessible after a block storage volume is deleted.

Uses snapshots to protect data against accidental or malicious deletion

Input/output operations per second (IOPS) and throughput are two performance metrics impacted by block size.

Block storage custom encryption services with VPC, such as IBM Cloud Hyper Protect Services and IBM Key Protect allow clients to utilize Keep Your Own Key (KYOK) and Bring Your Own Key (BYOK) capabilities.

Decouples data from user environments, allowing that data to be spread across multiple environments enabling the user to retrieve it quickly

Course 1: IBM Cloud Storage Options

File Storage

- Uses a hierarchical structure to organize files, folders, and subfolders
- Uses common protocols, such as NFS, SMBs, and CIFS
- Can be provisioned and then attached to one or multiple virtual machines or bare metal servers in the same zone or data center
- Provides strategic benefits such as reducing hardware costs and scaling up capacity
- Available in the endurance and custom tiers

Object Storage

- Manages unstructured data into self-contained units called objects
- Uses a flat structure known as a "bucket"
- Uses an application programming interface (API) to access and manage objects
- Stores encrypted and dispersed data across three geographic locations
- Offers different levels of resiliency: cross region, regional, and single data center
- Four tiers, including:
    - **Smart** – unknown and changing data patterns
    - **Standard** – data that is active and accessed frequently
    - **Vault** – less active data, typically accessed on a monthly basis
    - **Cold Vault** – data that is only accessed a few times a year
- Advantages to using object storage:
    - Enables clients to handle large amounts of unstructured data
    - Is scalable and cost-effective
    - Uses metadata allowing users to maximize the search feature and quickly access the object they need

Most internet communication, such as emails, photos, videos, audio files, sensor data, and other types of media and web content, is stored using object storage.

Course 1: IBM Cloud Storage Options

Software-Defined Storage (SDS)

- Is a storage architecture that separates storage software from its hardware
- Enables organizations to increase their storage capacity quickly making it flexible and scalable
- Allows the organization to use existing hardware that they currently have, which can be a tremendous cost savings
- Provides organizations the ability to store application data across different availability zones (if a zone goes down, all the compute and data will be available in another zone in the multizone region)

Red Hat® OpenShift® Data Foundation (ODF) and **Portworx** are two examples of SDS offerings used with IBM Cloud container orchestration services.

- ODF is a highly available storage solution for managing persistent storage for containerized apps. It consists of several open source technologies that allow users to provision and manage file, block, and object storage for containerized workloads in Red Hat OpenShift on IBM Cloud clusters.
- Portworx provides highly available persistent storage management for containerized apps. It does this by aggregating available storage that is attached to worker nodes and creating a unified persistent storage layer for containerized databases or other stateful apps that run in the cluster. Portworx comes with additional features, such as volume snapshots, volume encryption, isolation, and an integrated Storage Orchestrator for Kubernetes to ensure optimal placement of volumes in the cluster.

Question 1.

Data that is stored on block storage is secured on IBM Cloud by allowing clients to do what?

    A.   Clients can store their data in databases.

    B.   Clients can use Bring Your Own Key to encrypt their data.

    C.   Clients can move their data from one server to another.

    D.   Clients can perform processing or analysis on their data.

Answer B. Clients can use Bring Your Own Key to encrypt their data.

Question 2.

To access and manage objects when they are stored on IBM Cloud Object Storage, a client would use?

A. SQL

B. CIFS

C. API

D. SMB

Answer C. A client would use API to access and manage objects when they are stored on IBM Cloud Object Storage.

Question 3.

Describe how IBM Cloud Object Storage provides data resilience.

A. Data is always stored across at least three devices, zones, or regions.

B. IBM Cloud automatically backs the data up and restores it when needed.

C. The underlying database is replicated to a standby in another region.

D. Buckets are mirrored to underlying block storage in a second zone.

Answer A. IBM Cloud Object Storage provides data resilience by always storing data across at least three devices, zones, or regions.

Question 4.

What is the primary reason an organization using IBM Cloud would use software-defined storage?

- A. Being extremely fast, it is ideal for custom cloud native database services.
- B. It allows developers to trial disk layouts and IOPS configs where applications have demanding storage requirements.
- C. It has the ability to store application data across different availability zones, increasing resilience.
- D. It is the lowest cost of storage available, so it reduces total cost of ownership.

Answer C. The primary reason an organization using IBM Cloud would use software-defined storage is its ability to store application data across different availability zones, increasing resilience.

Question 5.

Which software-defined storage (SDS) option in IBM Cloud is highly available and can be used as persistent storage management for containerized applications?

A. Portworx

B. Contactor

C. NetApp Storage Appliance

D. DataCache

Answer A. Portworx in IBM Cloud is highly available and can be used as persistent storage management for containerized applications.

Question 6.

A client would like to migrate one of their workloads to IBM Cloud. They require low latency, high performance storage, and the storage needs to be highly redundant. Which IBM Cloud storage type meets the needs of this workload?

A. Block storage

B. Internal storage

C. File storage

D. Object storage

Answer A. Block storage would meet this client's workload needs.

Question 7.

The CIO for a media company requires media files stored in IBM Cloud to have high availability and the lowest cost. Which storage option would fulfill that requirement?

A.   File Storage

B.   NetApp Storage Appliance

C.   Block Storage for VPC

D.   Cloud Object Storage

Answer D. Cloud Object Storage would fulfill the need for a media company requiring media files stored in IBM Cloud to have high availability and the lowest cost.

Course 2: IBM Cloud Managed Databases

Introduction and Objectives

In Course 2 of the study guide, the subject matter:

- Compares and contrasts the different IBM Cloud-managed database options.

Lessons

- When to Use SQL and NoSQL Databases
- Benefits of Database-as-a-Service (DBaaS)
- IBM-Managed Cloud Databases
- Summary
- Knowledge Check

Objective

- Recognize the value of the different IBM Cloud-managed database options

Course 2: IBM Cloud Managed Databases

Two Categories of Databases

IBM provides many database offerings in IBM Cloud. Most of them fall into one of the following two categories:

Category	Description
Structured Query Language (SQL)	A SQL database supports a powerful language, known as SQL, which enables users to query, insert, update, and delete data in relational databases. Relational databases store data in fixed-structure tables. These tables are related to one another, for example, one table might contain employee names, another the address of each employee. The data is related using keys, which allow users to extract data by writing a SQL query that matches keys in each table.
Not only SQL (NoSQL)	NoSQL was developed in response to the emergence of cloud, big data, and web and mobile applications. Rooted in graph, document, key-value pairs, and wide-column stores, NoSQL adds horizontal scalability and loosens the rigidity of the table design that is the basis of SQL. NoSQL databases are non-relational. They store data in one data structure, such as a JSON document. Because non-relational designs don't require a schema, NoSQL databases can better manage large or unstructured data sets.

Course 2: IBM Cloud Managed Databases

Database-as-a-Service (DBaaS)

DBaaS is exactly what it sounds like: a service that lets users access and use a database over the cloud. With DBaaS, the provider handles the upgrades, backups, and other maintenance tasks to keep the system running 24/7, leaving the database administrator to focus on selecting the database that best fits with the workload and overall performance tuning. It is among the fastest-growing Platform-as-a-Service (PaaS) markets.

Benefits of DBaaS include:

- **Cost savings**: lets users pay a predictable, usage-based amount
- **Scalability:** lets users add storage and computing capacity when they need it and scale down during non-peak hours to reduce costs
- **Simplicity**: frees IT staff to concentrate on their primary roles
- **Rapid development**: lets developers request, configure, and integrate database capabilities in minutes
- **Security**: offers enterprise-grade security features such as default encryption of at-rest and in-transit data, integrated identity and access management controls, and measures to comply with regulatory standards
- **Reduced risk**: includes a service-level agreement (SLA) that guarantees a minimum level of uptime
- **Quality**: offers many configurable, preselected quality options that spare users the toil of wading through hundreds of databases

Course 2: IBM Cloud Managed Databases

IBM-Managed Database Options

IBM-managed SQL Databases

The following five IBM-managed SQL databases are enterprise ready, can be easily scaled and deployed, and provide encryption and high availability options:
- Db2® on Cloud
- IBM Cloud Databases for Postgres
- IBM Cloud Databases for Enterprise DB
- IBM Cloud Databases for MySQL
- IBM Db2 Warehouse on Cloud

IBM-managed NoSQL Options

The following six IBM-managed NoSQL options offer global availability, are easy to scale, easy to deploy, and have multiple encryption and high availability options:
- Cloudant®
- IBM Cloud Databases for DataStax
- IBM Cloud Databases for Elasticsearch
- IBM Cloud Databases for MongoDB
- IBM Cloud Databases for Redis
- IBM Messages for RabbitMQ

Question 1.

One type of database IBM Cloud offers is SQL. What is another type of database IBM Cloud offers?

A. NoSQL
B. Persistent
C. Shared
D. Private

Answer A. SQL and NoSQL are two types of databases IBM Cloud offers.

Question 2.

The _____ type of database was developed in response to the emergence of cloud, big data, web, and mobile applications.

A. SQL

B. NoSQL

C. Persistent

D. Shared

Answer B. NoSQL was developed in response to the emergence of cloud, big data, web, and mobile applications.

Question 3.

Who is responsible for handling upgrades, backups, and other maintenance tasks with IBM Cloud's Database-as-a-Service offerings?

A.   Database vendor

B.   Shared responsibility

C.   Application owner

D.   Service provider

Answer D. The service provider is responsible for handling upgrades, backups, and other maintenance tasks with IBM Cloud's Database-as-a-Service offerings.

Course 3: Storage Security Options

Introduction and Objectives

In Course 3 of the study guide, the subject matter:

- Explains how role-based access control (RBAC) is used to secure solutions built on IBM Cloud, the encryption options available on IBM Cloud, including the encryption options for storage when utilizing IBM Cloud VMware Solutions.

Lessons

- RBAC Secures Solutions Built on IBM Cloud
- Encryption Options Available in IBM Cloud
- Secure Storage Options for IBM Cloud VMware Solutions
- Summary
- Knowledge Check

Objectives

- Identify storage security options for IBM Cloud

Course 3: Storage Security Options

Role-Based Access Control (RBAC)

RBAC is one method used to enhance security on IBM Cloud. The following RBAC methods secure solutions on IBM Cloud:

User Identity and Access Management (IAM)

It is important to only allow access to services to those who need it and to revoke access when it is no longer needed. IAM allows administrators to grant privileges to individual users or to grant and control access to large numbers of users via **Access Groups.**

OpenShift & Containers

In a container environment, cluster administrators can use the cluster roles and bindings to control who has various access levels to the OpenShift Container platform itself and all projects. Developers can use local roles and binding to control who has access to their projects.

VMware

All users that access IBM Cloud for VMware Solutions services must be assigned an access policy. This access policy must be paired with a defined IAM user role. The assigned access policy dictates the actions the user can complete.

Course 3: Storage Security Options

Data Encryption

Data can be protected by encrypting it. This is done using keys. Once the data is encrypted, unless you have the encryption key, decrypting the data so that it can be read is extremely difficult. In the cloud, there are **three places where data is encrypted**.

Data is at rest when it's in storage. Strategies such as disk encryption protect data at rest.

Data is in motion when it's being transferred: storage to processing, processing to storage, or to and from the end user. Strategies such as using encrypted transmission protocols like Transport Layer Security (TLS) protect data when it's in motion.

Data in use by the central processing unit (CPU) and memory is the most vulnerable to attacks.

Encryption Types

The key to encryption is the cipher key. There are two primary ways to manage storage of the cipher key:
1. Using supplier-managed encryption
2. User client-managed encryption

IBM-Managed Encryption

When provisioning compute resources, IBM automatically applies encryption to their attached disk volumes to protect data. There are two volumes that are encrypted by default:
1. Primary boot volumes
2. Secondary boot volumes

Client-Managed Encryption

Client-managed encryption enables users to bring their own client root key (CRK) or use a key management service (KMS) to generate a key. KMSs that are provided by IBM Cloud are IBM Key Protect for IBM Cloud and IBM Hyper Protect Crypto Services (HPCS). In addition, there are several KMSs specific to VMware that clients can also choose from.

Course 3: Storage Security Options

IBM Key Protect for IBM Cloud

IBM Key Protect keeps data secure using a process called envelope encryption, where one key is encrypted with another key. For example, the root key encrypts the data encryption keys (DEKs) that encrypt the plaintext data. Root keys protect the DEKs and are managed by IBM in an impenetrable hardware security module (HSM).

When decrypting the data, the encrypted DEK is opened first (opening the envelope) using the root key, and then the DEK is used to decrypt the actual data.

Key Management Interoperability Protocol (KMIP)

KMIP is an encryption service specific to VMware that can be used with IBM Cloud. Using KMIP for VMware service, alongside IBM Cloud Key Protect or IBM Cloud HPCS to enable vSAN or VMware vSphere® provides encryption for workloads.

There are four VMware solution platforms that offer security options:
• Client-managed platforms :
  • vCenter® Server
  • vSphere® Server
  Both vCenter Server and vSphere Server are single-tenant models with higher levels of isolation for enhanced security. Both have two encryption capabilities available: encryption in transit and encryption at rest.
• IBM-managed platforms:
  • VMware as a Service
  • VMware Shared
  Both VMware as a Service and VMware Shared provide standard and custom deployment options.

Question 1.

_____ can be used to encrypt data in motion.

A. File Transport Protocol (FTP)

B. Transport Layer Security (TLS)

C. Hyper Text Transport Protocol (HTTP)

D. Transport Security Layer (TSL)

Answer B. TLS can be used to encrypt data in motion.

Question 2.

Data at rest is where data is stored physically. Data is motion is:

    A.   Data moving from one place to another

    B.   Data in an operational state

    C.   Data being retrieved from storage

    D.   Data used over mobile devices

Answer A. Data at rest is where data is stored physically, and data in motion is data moving from one place to another.

Question 3.

Key Protect is a key management solution on IBM Cloud. What is another key management solution on IBM Cloud?

A.   Block Storage

B.   Object Storage

C.   Identity and Access Management

D.   Hyper Protect Crypto Services

Answer D. Key Protect and Hyper Protect Crypto Services are two key management solutions on IBM Cloud.

Course 4: IBM Hyper Protect Crypto Services

Introduction and Objectives

In Course 4 of the study guide, the subject matter:

- Reviews the benefits and options of IBM Hyper Protect Crypto Services (HPCS).

Lessons

- IBM Hyper Protect Crypto Services Features
- Key Management Concepts
- Integration with IBM Hyper Protect Crypto Services
- Use Cases for IBM Hyper Protect Crypto Services
- Summary
- Knowledge Check

Objectives

- List the benefits and options of IBM Hyper Protect Crypto Services

Course 4: IBM Hyper Protect Crypto Services

Hyper Protect Crypto Services (HPCS)

IBM offers an enhanced level of client-managed encryption, FIPS 140-2 Level 4 certified, referred to as Keep Your Own Key (KYOK) solution. This KYOK solution is available for clients who want the highest level of data protection. The level is so high that even IBM Cloud has no access to their keys. Only the client can decrypt the data stored in IBM Cloud. This service is called IBM Cloud Hyper Protect Crypto Services (HPCS).

HPCS is a dedicated key management service that uses a dedicated, client-controlled, cryptographic processing hardware security module (HSM) to generate, encrypt, store, and decrypt keys.

A single-tenant key management service with key vaulting, provided by the HSMs helps users easily create and manage their encryption keys. They can either create encryption keys using HPCS or bring their own keys when they set up the service. They also retain exclusive control of their keys with the KYOK feature.

Key Differences between IBM Key Protect and HPCS

	Differences
IBM Key Protect	• Includes FIPS 140-2 Level 3 certification and is often referred to as a Bring Your Own Key (BYOK) solution • The cloud provider has already performed these tasks, which means they are the true owners of the root of trust. • Is a multi-tenant key-management service • Uses a shared HSM
HPCS	• Includes FIPS 140-2 Level 4 certification and is often referred to as Keep Your Own Key (KYOK) • For ultimate protection, users must control and initialize the HSM. This is only possible with KYOK services like HPCS. • Is a single-tenant key-management service • Uses a dedicated HSM

Question 1.

What is one of the differences between IBM Key Protect and IBM Hyper Protect Crypto Services?

A. IBM Key Protect service supports the Keep Your Own Key capability, and IBM Hyper Protect Crypto Services supports Bring Your Own Key capability.

B. IBM Key Protect is a shared multi-tenant service, and IBM Hyper Protect Crypto Services is a single-tenant service.

C. IBM Key Protect and IBM Hyper Protect Crypto Services support the same features and are only comparable on price.

D. IBM Key Protect is a single-tenant service, and IBM Hyper Protect Crypto Services is a shared multi-tenant service.

Answer B. IBM Key Protect is a shared multi-tenant service, and IBM Hyper Protect Crypto Service is a single-tenant service.

Question 2.

A client needs a key management solution on IBM Cloud that meets FIPS 140-2 Level 4 certification. Which option should be recommended?

A. IBM Cloud HSM

B. IBM Blockchain Platform

C. IBM Hyper Protect Crypto Services

D. IBM Key Protect

Answer C. IBM Hyper Protect Crypto Services meets the FIPS 140-2 Level 4 certification.

Question 3.

IBM Hyper Protect Crypto Services provides the highest available encryption certification. Which level is the highest?

A. SEC 17a-f

B. FIPS 140-2 Level 3

C. FIPS 140-2 Level 4

D. FedRAMP Maximum

Answer C. The highest level of available encryption certification is FIPS 140-2 Level 4.

Question 4.

One of the primary differences between IBM Key Protect and IBM Hyper Protect Crypto Services (HPCS) is:

A.   HPCS is only available in the client's data center, and Key Protect is available in IBM Cloud.

B.   HPCS cannot crypto shred data, and Key Protect can crypto shred data.

C.   HPCS is for encrypting mainframe data only, and Key Protect is for x86 based platform data.

D.   HPCS uses a dedicated HSM, and Key Protect uses a shared HSM.

Answer D. HPCS uses a dedicated HSM, and Key Protect uses a shared HSM.

Question 5.

A financial client needs encryption for their environment. They would like to be the only one with the ability to access their data on their production environment. Which service in IBM Cloud would be available for this client to use that meets their needs?

A. Hyper Protect Crypto Services

B. Security and Compliance Center

C. Identity and Access Manager

D. Key Protect

Answer A. Hyper Protect Crypto Services would meet this client's requirements.

Course 5: IBM Backup, Recovery, and Replication

Introduction and Objectives

In Course 5 of the study guide, the subject matter:

- Explores IBM Cloud backup, recovery, and restore options, their capabilities, and benefits.

Lessons

- Backup and Recovery Capabilities
- Backup and Disaster Recovery for x86 Workloads
- Backup and Disaster Recovery for Power Systems Virtual Servers
- Summary
- Knowledge Check

Objectives

- Recognize the benefits and options of IBM Backup, Recovery, and Replication Solutions

Course 5: IBM Backup, Recovery, and Replication

Backup and Restore
Backup and restore make periodic copies of data and applications to a secondary device. Clients can then use those copies to recover the data in the event of an unplanned event, such as a power outage, cyberattack, human error, or a disaster.

Disaster Recovery
Disaster recovery is a plan that is created to deal with an outage that impacts application, data, and IT resources.

When working with a client, recovery time objective (RTO) and recovery point objective (RPO) should be discussed.

Recovery Time Objective (RTO)

The maximum amount of time the organization can afford to be without access to the data or application.

Recovery Point Objective (RPO)

The amount of data an organization can afford to lose and effectively dictates how frequently they need to back up their data to avoid losing more.

Course 5: IBM Backup, Recovery, and Replication

Disaster Recovery for x86 Workloads

There are a variety of services available for disaster recovery, and IBM uses and offers services such as:

IBM Cloud Backup
- Service that provides a method to back up data between servers in one or more data centers on the IBM Cloud network
- Can be set for a daily, weekly, or custom schedule
- Can target the entire system, a specific directory, or an individual file

Zerto
- Service that provides replication and disaster recovery capabilities
- Can be integrated into the deployment offerings to protect and recover data in a VMware virtual environment on IBM Cloud
- Uses continuous data replication with journaling versus snapshots

Veeam
- Service that integrates directly with VMware hypervisors
- Helps an organization achieve high availability and resiliency to control backup and restore for all the virtual machines attached to their infrastructure from one console

IBM Spectrum Protect Plus
- Creates an efficient and scalable solution for clients who need data protection, reuse, and recovery in virtual environments
- Can be implemented as a stand-alone solution or part of data governance and long-term storage solution with IBM Spectrum Protect service

Course 5: IBM Backup, Recovery, and Replication

Disaster Recovery for IBM Power Systems Virtual Servers

In the event of a hardware failure, the IBM Power Systems Virtual Server service will restart the virtual servers on a different host system to provide uninterrupted service. This process provides basic High Availability (HA) capabilities for Power Systems Virtual Servers.

There are three different services available on IBM Cloud to back up Power System Virtual Servers based on operating system:
1. Image Capture: utilizes a FlashCopy of the logical partition (LPAR) to recover data
2. AIX Backups: utilizes strategies like Veeam and IBM Spectrum Protect
3. IBM i Backups: Uses IBM Cloud Storage Solutions (ICC) along with IBM Backup Recovery and Media Services (BRMS) to automatically back up LPARS (logical partitions) to Cloud Object Storage

The following table indicates which backup option is available for each Power Systems Virtual Server operating system.

	Veeam	IBM Spectrum Protect	IBM Cloud Storage Solutions (ICC)	FalconStor StoreSafe VTL
AIX	X	X		
IBM i			X	X

Course 5: IBM Backup, Recovery, and Replication

Replication Options for Power Systems Virtual Servers

Multiple disaster recovery options are available for Power Systems Virtual Servers. These options use some form of replication, whether it be storage replication, operating system replication, or logical replication solutions from third-party vendors.

There are four replication options:

Storage Replication using Global Replication Services (GRS)

GRS maintains a consistent and recoverable copy of the data at a remote site. This option may only be used between two Power System Virtual Server locations.

Host or Operating System Mirroring

IBM® PowerHA® technology is available for Power Systems Virtual Servers with both IBM AIX and IBM i operating systems.

The PowerHA SystemMirror® software monitors the system, and when a component fails, it automatically moves the application (along with the resources that ensure access to the application) to another node in the cluster, ensuring continued access.

For AIX operating systems, a disaster recovery mechanism between two AIX virtual server instances in separate IBM Cloud data centers can be implemented using geographic logical volume manager (GLVM) replication. GLVM replication and PowerHA software provide clients with a fully automated solution.

For IBM i operating systems, a disaster recovery mechanism between two IBM i virtual server instances in separate IBM Cloud data centers can be implemented using PowerHA geographic mirroring.

Application and Database Level Replication

Clients may select to replicate selected applications or databases.

Logical Replication Solutions

Additional replication solutions are available from various vendors, including Mimix, Rocket, Carbonite, Maxima, and QuickEDD. These solutions use a variety of methods to assist clients in preventing data loss and unplanned downtime.

Question 1.

Disaster recovery is defined as:

A. A plan that is created to deal with an outage that could impact applications, data, and IT resources

B. A plan that is created to determine how long a copy of data will be stored in a secondary location

C. A plan that is created to determine how often data will be replicated to a secondary region

D. The process of making periodic copies of data and applications to a secondary device

Answer A. Disaster recover is defined as a plan that is created to deal with an outage that could impact applications, data, and IT resources.

Question 2.

Recovery time objective is defined as:

A. A plan that is created to deal with an outage that could impact applications, data, and IT resources

B. The device a client needs to use to backup their data and applications prior to an incident

C. The amount of data an organization can afford to lose during an incident

D. The maximum period in which the organization can afford to be without access to data during an incident

Answer D. Recovery time objective is defined as the maximum period in which the organization can afford to be without access to data during an incident.

Question 3.

When running VMware Solutions Dedicated, which disaster recovery and protection option on IBM Cloud uses continuous data replication with journaling versus snapshots?

    A.   IBM Spectrum Protect Plus

    B.   CommVault

    C.   Zerto

    D.   IBM Cloud Backup

Answer C. Zerto is the disaster recovery and protection option on IBM Cloud that uses continuous data replication with journaling.

Question 4.

An organization is leveraging IBM Cloud Power Systems Virtual Servers for their AIX workloads. Veeam is one solution that could be used for file-based backup and disaster recovery for this environment. What is another solution that could be used?

A.   FalconStor StorSafe VTL

B.   IBM Cloud Spectrum Protect

C.   Zerto

D.   Storage Solutions

Answer B. IBM Cloud Spectrum Protect is another solution that could be used.

Question 5.

A client is leveraging IBM Cloud Power Systems Virtual Servers for their IBM i workloads. The _____ disaster recovery mechanism could be used.

A. PowerHA Geographic Mirroring

B. Veeam

C. IBM Spectrum Protect

D. Geographic Logical Volume Manager

Answer A. PowerHA Geographic Mirroring could be used for a client who is leveraging IBM Cloud Power Systems Virtual Servers for their IBM i workloads.

Acronyms

Acronym	Acronym Expansion
ACID	Atomicity, Consistency, Isolation, and Durability
API	Application Programming Interface
BYOK	Bring Your Own Key
CIFS	Common Internet File System
CLI	Command Line Interface
CRK	Client Root Key
DBaaS	Database as a Service
DR	Disaster Recovery
DEK	Data Encryption Key
FISMA	Federal Information Security Management Act
FIPS	Federal Information Processing Standard
GRS	Global Replication Service
HDD	Hard Disk Drive
HA	High Availability
HIPAA	Health Insurance Portability and Accountability Act
ICC	IBM Cloud Storage Solution
KMIP	Key Management Interoperability Protocol
KMS	Key Management Service
HPCS	Hyper Protect Crypto Services
HSM	Hardware Security Module

Acronyms

Acronym	Acronym Expansion
IAM	Identity and Access Management
IOPS	Input/Output Operations Per Second
IoT	Internet of Things
KMIP	Key Management Interoperability Protocol
MZR	Multizone Region
NAS	Network Attached Storage
NFS	Network File System
NoSQL	Not Only SQL
ODF	OpenShift Data Foundation
PCI	Payment Card Industry
PKCS	Public-Key Cryptography Standards
RBAC	Role-Based Access Control
RPO	Recovery Point Object
RTO	Recovery Time Object
SaaS	Software as a Service
SANs	Storage Area Networks
IAM	Identity and Access Management
IOPS	Input/Output Operations Per Second
IoT	Internet of Things
KMIP	Key Management Interoperability Protocol

Acronyms

Acronym	Acronym Expansion
SDS	Software-Defined Storage
SLA	Service-Level Agreement
SMB	Server Message Block
SQL	Structured Query Language
SSD	Solid-State Drive
TDE	Transparent Data Encryption
VM	Virtual Machine
VMFS	Virtual Machine File System
VPC	Virtual Private Cloud
VSI	Virtual Server Instance
WORM	Write-Once-Read-Many

Courses and Objectives

Courses

1. IBM Cloud Integration Use Cases
2. Edge Solutions on Cloud
3. IBM Cloud AI and Analytics Use Cases

Objectives

- Identify integration options for IBM Cloud® and list the benefits they provide
- Identify options for IBM Cloud Edge Solutions and list the benefits they provide
- Identify options for IBM Cloud AI and analytics and list the benefits they provide

Course 1: Cloud Integration Use Cases

Introduction and Objectives

In Course 1 of the study guide, the subject matter:

▪ Explores use cases for integration on IBM Cloud, including IBM API Connect®, APP Connect, IBM Messaging Options, IBM Event Streams, and IBM MQ.

Lessons

▪ API Connect, APP Connect, and IBM Messaging Options

▪ Summary

▪ Knowledge Check

Objective

• Identify integration options for IBM Cloud and list the benefits they provide

Course 1: Cloud Integration Use Cases

Why integrate?
Software-as-a-Service (SaaS) applications have become common, and artificial intelligence (AI) capabilities require organizations to integrate broad ranges of data from rapidly multiplying sources. Integration has become increasingly necessary.

Some organizations have adapted to the added difficulty of integration by replacing large code silos with small, independent components called *microservices*. This strategy offers greater scalability, resilience, and agility, which is why it is called **agile integration.**

Agile integration decentralizes services, so application teams can better control the creation and exposure of their application program interfaces (APIs), messages, and events. APIs are critical for bringing the correct information together at the correct time, in the correct context. IBM offers several technologies and tools to integrate APIs, events, and messages.

IBM API Connect
Allows users to create, secure, expose, manage, socialize, and analyze its APIs across clouds.

IBM App Connect
Integrates data and applications from existing systems and ties together technologies across environments, including legacy and SaaS systems.

IBM Event Streams
Apache Kafka-based IBM Event Streams enable organizations to create smart applications that react to events as soon as they happen.

IBM MQ
Scalable messaging platform that uses Java Message Service (JMS) and other technologies to integrate applications and encrypt the data end-to-end.

IBM App Connect Enterprise
Helps integration specialists develop, deploy, and support message flow applications by connecting applications together, regardless of the message formats or protocols they support.

Course 1: Cloud Integration Use Cases

Which integration solution should be used?

If you need to:	Consider using:
Create, secure, and manage APIs	IBM API Connect
Connect, copy, and sync application data across platforms	IBM App Connect
Deliver reliable, secure messaging, either on-premises or over the cloud	IBM MQ
Build adaptive smart apps that react to events as soon as they happen	IBM Event Streams

Question 1.

The _____ service lets users create, secure, expose, manage, socialize, and analyze APIs across clouds.

A. IBM Event Streams

B. IBM App Connect Enterprise

C. IBM API Connect

D. Simple Storage Service (S3) API

Answer C. IBM API Connect lets users create, secure, expose, manage, socialize, and analyze APIs across clouds.

Question 2.

Which IBM Cloud service would be appropriate to recommend to a client who would like to use Apache Kafka?

    A.   IBM API Connect

    B.   IBM App Connect Enterprise

    C.   IBM Event Streams

    D.   IBM Cloud Pak for Security

Answer C. IBM Event Streams would be appropriate to recommend to a client who would like to use Apache Kafka.

Course 2: Edge Solutions on Cloud

Introduction and Objectives

In Course 2 of the study guide, the subject matter:

▪ Explores IBM Edge Computing® and how Edge computing might leverage IBM Cloud

Lessons

▪ How Edge Functions Work
▪ IBM Cloud Satellite
▪ Summary
▪ Knowledge Check

Objective

▪ Identify options for IBM Cloud Edge Solutions and list the benefits they provide

Course 2: Edge Solutions on Cloud

What is edge computing?

Edge computing is a distributed computing topology that aims to use bandwidth more efficiently, by bringing applications close to where data is created, and actions are performed.

Developers can use a serverless execution environment to create and modify existing applications without having to configure or maintain infrastructure. By eliminating the need to send data over a network for processing, edge computing reduces latency. This lets users analyze the data faster and more thoroughly, extract deeper insights, and respond to requests more quickly for a better client experience.

IBM Cloud Internet Services

One way IBM Cloud offers edge solutions is through IBM Cloud Internet Services (CIS). CIS plays a significant part in assisting global load balancers distribute workloads across a network. Because CIS has so many data centers around the world, it can typically supply compute power very close to the location of the data.

See below how an IBM CIS Edge function delivers serverless computing capability to users across the globe.

Assign URIs to Actions

When implementing a CIS Edge function, developers consider the actions they want it to perform. For each action, they define uniform resource identifiers (URIs) in the domain. The associations between actions and URIs are called triggers.

Check for a Match

When a user sends an HTTP request, IBM CIS Edge intercepts it at the Cloud edge and reviews the URL for matches with any triggers in the domain.

Run the Action

If the edge function finds a match, a trigger event causes IBM CIS Edge to run the associated action.

Course 2: Edge Solutions on Cloud

IBM Cloud Internet Services (Continued)

- In addition to handling HTTP requests and responses, IBM CIS Edge supports functions that can make parallel requests or generate responses from the cloud edge.
- If processing needs to occur on client-provided servers physically located at the edge, IBM Cloud Satellite® might be a more appropriate solution.
- Because edge functions run code on trusted CIS Edge servers, users don't need to use a modern browser.

IBM Cloud Satellite is IBM's distributed cloud offering, which provides consistency in services and the ability to deploy anywhere. It is a platform service that enables organizations to run IBM Cloud services in the location of their choice: on premises, a competitor's cloud, or in edge networks (distributed cloud framework that brings enterprise applications closer to a data source).

Benefits of IBM Cloud Satellite

✓ Clients can deploy IBM Cloud managed services in the supported location of their choice, enabling them to run services and applications close to other workloads or to adhere to security requirements.

✓ Clients provide their own infrastructure. It can reside in their on-premises data center, in another cloud, or at an edge location. IBM manages the Satellite installation, while clients maintain the physical components.

✓ Clients can provision Satellite-ready apps into their location, including Red Hat® OpenShift® and IBM Cloud Databases, directly from the IBM Cloud Catalog.

✓ Satellite locations are securely connected to IBM Cloud by the means of an encrypted transport layer security (TLS) tunnel.

✓ Clients have access to IBM Cloud Monitoring, IBM Log Analysis, and IBM Cloud Activity Tracker within the IBM Cloud Satellite.

Course 2: Edge Solutions on Cloud

Main Satellite Components

Satellite Hosts

- Minimum of three hosts are required for a Satellite location
- Client-provided computer running Red Hat® Enterprise Linux 7.x or higher

Satellite Link

- Encrypted TLS tunnel that allows clients to securely connect the Satellite installation to other services
- Can route requests for multiple destination resources using transmission control protocol (TCP) ports 80 (HTTP) and 443 (HTTPS)

Satellite Location

- Where the server and storage infrastructure resides, onto which Satellite-enabled services are deployed

IBM Cloud Object Storage

- An instance of IBM Cloud Object Storage is needed to securely store information about the Satellite configuration.

Endpoints

- Satellite uses two types of endpoints: location endpoint and cloud endpoint.
- A location endpoint allows resources in IBM Cloud to securely access a resource in an on-premises Satellite location.
- A cloud endpoint allows resources in on-premises Satellite locations to access a resource that runs anywhere outside of the Satellite location.

Question 1.

Hosts, links, location, and _____ are the four main components of an IBM Cloud Satellite solution.

A.   Endpoints

B.   API

C.   Direct Link

D.   Client server

Answer A. Hosts, links, location, and endpoints are the four main components of an IBM Cloud Satellite solution.

Question 2.

Edge functions run in an IBM Cloud CIS Edge solution on _____ .

A.   Servers in IBM Cloud multizone region (MZR) data centers

B.   IBM CIS Edge servers in one of many global CIS data centers

C.   An IBM Storage Fusion appliance at the edge

D.   An IBM Cloud Satellite location for low latency

Answer B. Edge functions run in an IBM Cloud CIS Edge solution on IBM CIS Edge servers in one of many global CIS data centers.

Question 3.

The IBM Cloud Satellite component, _____ , uses the internet to provide a secure connection between the satellite location and other IBM Cloud services, such as monitoring and management.

A.   Security Transport Service

B.   Direct Link

C.   Satellite Link

D.   IBM Cloud Virtual Private Network

Answer C. The IBM Cloud Satellite component, Satellite Link, uses the internet to provide a secure connection between the satellite location and other IBM Cloud services, such as monitoring and management.

Question 4.

A client has a division in a country with strict data residency requirements. Which IBM Cloud service would guarantee data stays within the borders of the country and would allow consumption of cloud services?

A.   Internet Services Edge

B.   Satellite

C.   Virtual Private Cloud

D.   Content Delivery Network

Answer B. Satellite would allow consumption of cloud services and guarantee data stays within the borders of the country.

Question 5.

IBM Cloud _____ securely stores information about an IBM Cloud Satellite configuration.

A.   Wasabi Hot Cloud Storage

B.   Block Storage Snapshots for VPC

C.   Object Storage

D.   Cloud Backup for VPC

Answer C. IBM Cloud Object Storage securely stores information about an IBM Cloud Satellite configuration.

Question 6.

Organizations have the option to run select IBM Cloud Services on premises, behind their corporate firewall, at the edge, or in competitors' clouds when using which IBM Cloud service?

A.   IBM Storage Fusion

B.   Red Hat Satellite

C.   IBM Cloud Pak for Multi Cloud Management

D.   IBM Cloud Satellite

Answer D. Organizations have the option to run select IBM Cloud Services on premises, behind their corporate firewall, at the edge, or in competitors' clouds when using IBM Cloud Satellite.

Course 3: IBM Cloud AI and Analytics Use Cases

Introduction and Objectives

In Course 3 of the study guide, the subject matter:

- Explores artificial intelligence (AI) and analytics use cases on IBM Cloud.

Lessons

- IBM's AI Ladder
- IBM Services that Support Analytics and AI
- Use Cases Associated with Analytics and AI
- Summary
- Knowledge Check

Objective

Identify options for IBM Cloud AI and analytics and list the benefits they provide

Course 3: IBM Cloud AI and Analytics Use Cases

IBM thinks of the process of implementing AI in terms of a ladder with four rungs.
The AI Ladder allows companies to simplify and automate the process of converting data into useful information.
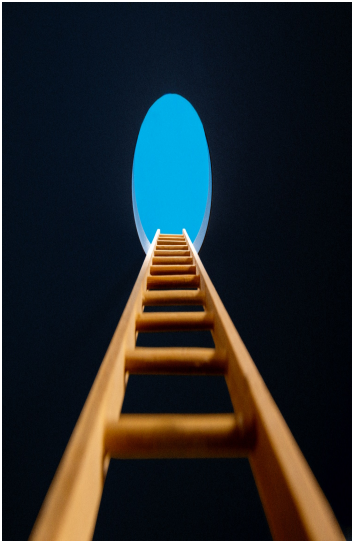
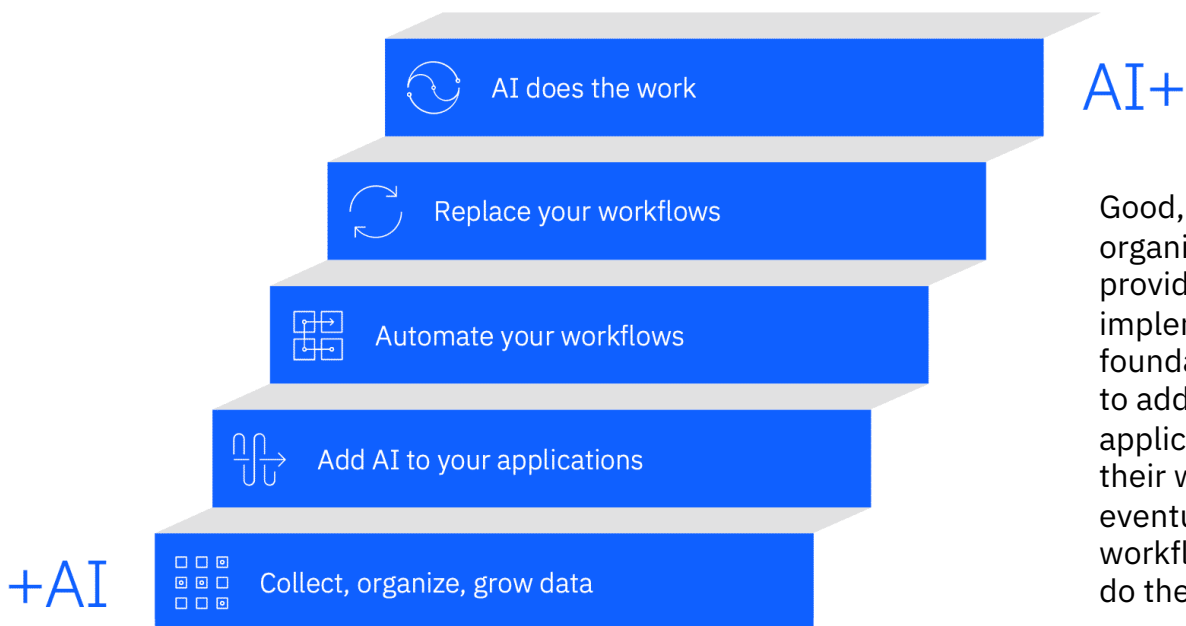**Infuse** – Apply AI models to enhance organizational workflows through advanced analytical tools

Analyze – Build AI models to scale and manage business insights everywhere

Organize – Cataloging data and adding governance to create a trusted analytic foundation to make data meaningful to an organization

Collect – Make data easy to consume and access

Modern-Day AI Ladder

AI has shifted from a world where companies are thinking about running their business plus including some AI to a world that is primarily focused on AI to help the business. The way organizations think of AI has shifted: +AI to AI+

AI+

AI does the work

Replace your workflows

Automate your workflows

Add AI to your applications

+AI

Collect, organize, grow data

Good, collected, well-organized, and trusted data provides the foundation for implementing AI. This foundation allows clients to add AI to their applications and automate their workflows and eventually, to replace their workflows and allow AI to do the work.

Course 3: IBM Cloud AI and Analytics Use Cases

IBM AI Services

IBM offers many services that use AI and analytics to improve business processes, including Watson® Assistant, Watson Studio, and Watson OpenScale.

Watson Assistant	Watson Studio	Watson OpenScale
• Builds dynamic, branded conversational interfaces into any application, device, or channel • Gives customers fast, accurate, and consistent information • IBM Watson Assistant for Citizens brings together Watson Assistant's Natural Language Processing capabilities with Watson Discovery and uses client-provided information to automate responses to frequently asked questions. It's available in English and 14 other languages.	• Embeds AI into a business by creating flexible, custom data models • Tools include Jupyter Notebook, an open-source web application for creating/sharing documents with live code, equations, visualizations, and text and RStudio®, which lets users share data across organizations.	• Tracks outcomes through the AI life cycle (scope, design/build, and deploy) • Monitors risk for performance, bias, and explainability • Creates fair and explainable outcomes and measures how models and business goals align • IBM AI Explainability 360 is an OpenScale-powered software toolkit that identifies where explainability is needed and provides a platform to incorporate them.

Question 1.

The IBM Cloud service that collects user input to help train consistent responses to user's questions across any application, device, or channel is called _____ .

A.   Watson Assistant

B.   Natural Language Processing

C.   Watson Tone Analyzer

D.   Watson OpenScale

Answer A. The IBM Cloud service that collects user input to help train consistent responses to user's questions across any application, device, or channel is called Watson Assistant.

Question 2.

_____ monitors risk models for performance, bias, and explainability to create more fair and explainable outcomes for clients.

A. Watson OpenScale

B. Watson Tone Analyzer

C. Natural Language Processing

D. Watson Assistant

Answer A. Watson OpenScale monitors risk models for performance, bias, and explainability to create more fair and explainable outcomes for clients.

Question 3.

A client would like to deploy an online tool to help answer their employees' questions in real-time based on its bank of frequently asked questions (FAQs) and natural language processing to cover all the supported languages. Which IBM Cloud service would help this client?

A. Watson Tone Analyzer

B. Watson OpenScale

C. Watson Assistant

D. AI Explainability 360

C. Watson Assistant would help a client who would like to deploy an online tool to help answer their employees' questions in real-time based on its bank of FAQs and natural language processing to cover all supported languages.

Acronyms

Acronym	Acronym Expansion
AI	Artificial Intelligence
API	Application Programming Interface
CIS	Cloud Internet Services
JMS	Java Message Service
SaaS	Software as a Service
TCP	Transmission Control Protocol
TLS	Transport Layer Security
URI	Uniform Resource Identifiers

Courses and Objectives

Courses

1. Building Cloud Native Applications
2. Modernizing Existing Applications

Objectives

- Identify the key enabling technologies of cloud native applications and explain the benefits they provide
- Describe the benefits of modernizing existing applications

Course 1: Building Cloud Native Applications

Introduction and Objectives

In Course 1 of the study guide, the subject matter:

- Explores the core principles and practices of building cloud native applications.

- Discusses microservices and their benefits and the key tools and technologies used in applying cloud native solutions.

Lessons

- Cloud Native Benefits and Use Cases
- Twelve-Factor App Methodology
- Container Orchestration and Serverless Architectures
- REST APIs and API Gateways
- Messaging and Event Streaming
- DevSecOps
- Authentication
- Infrastructure as Code
- Summary
- Knowledge Check

Objectives

- Identify the key enabling technologies of cloud native applications and explain the benefits they provide

Course 1: Building Cloud Native Applications

Cloud native is software that is designed to run in the cloud. Instead of traditional applications that are installed locally, cloud native applications are designed to run exclusively in public or private clouds via the internet.

There are advantages and disadvantages to cloud native applications:

Advantages	Disadvantages
• Creating/updating an application is easier (composed of multiple, smaller microservices rather than one monolithic application). • Scaling is targeted - only need to scale the microservices that need it. • Updates made without downtime. • Microservices can be used with multiple applications, bringing benefits of consistency.	• Instead of one large application, it requires several small, discrete services. • Different toolsets are needed since cloud native applications use a DevSecOps pipeline to continuously deploy changes. • Corporate culture of the organization needs to be able to handle the paradigm shift to rapid development.

Course 1: Building Cloud Native Applications

Microservice Architecture
Instead of building applications where all the component parts are contained within one monolithic application, microservices architecture breaks up applications into a conglomeration of loosely coupled, smaller services.
Features:
- Independently deployable
- Specialized functionality
- Easy scaling

The **Twelve-Factor App methodology** was created in 2012 as a set of guidelines to assist developers in creating cloud native applications.

The goals of this methodology were to design cloud native applications that:
➢ Use a declarative approach, reducing the time needed to onboard new developers.
➢ Have an organized, reliable interface with the underlying operating system.
➢ Can be deployed on modern cloud platforms.
➢ Maximize the connection between development and production.
➢ Have the ability to scale up quickly without major changes to architecture or development processes.

There are 12 factors in this methodology: Codebase; dependencies; configuration; backing services; build, run, release; processes; port binding; concurrency; disposability; dev/prod parity; logs; and admin processes.

Course 1: Building Cloud Native Applications

Key Enabling Technologies for a Cloud Native Solution

Container Orchestration	• Containers are self-contained applications that can operate on their own. Instead of being included in only one application, microservices can be re-used in multiple applications. • Microservices are deployed into containers. A container holds all the executable code, libraries, and other dependencies that the microservice requires to execute. • **Kubernetes** (sometimes shortened to K8s) is used to deploy and manage containers across a cluster of host servers. It controls the resources that each container uses and ensures that the container is running. • The primary container orchestration services provided on IBM Cloud are IBM Kubernetes Services and Red Hat® OpenShift®.

Serverless Architecture	• Serverless is a cloud computing executable model that is designed to eliminate the common issues with servers.
	• Features include:
	   • Resources required to run application code on demand are automatically provisioned.
	   • Resources are automatically scaled up or down depending on demand.
	   • Resources are automatically scaled to zero when the application stops running.
	• Although servers are used in a serverless model, they are transparent to the end user. All provisioning, scaling, patching, and other maintenance tasks are the responsibility of the cloud provider.

Course 1: Building Cloud Native Applications

Key Enabling Technologies for a Cloud Native Solution (Continued)

REST APIs and API Gateways	• An application programming interface (API) is a set of rules that allow applications or devices to communicate with each other. • They are used to communicate data and commands between microservices. • They provide security and authentication as well as transmit and receive commands and data. • Representational State Transfer (REST) APIs have additional architectural constraints to them.
Messaging and Event Streaming	Messaging and event streaming are designed to broadcast state changes to other services such as error logging services. Since microservice instances are independent, maintaining an error log for each instance is impractical. Using a centralized service that collects and analyzes error data is the best technique.
DevSecOps/Continuous Delivery	DevSecOps is a software development approach that combines development (Dev), security (Sec), and operations (Ops) practices into a unified process. Continuous delivery is part of the overall DevSecOps process. Continuous delivery is the process of preparing software for production in short cycles and automating the process that moves code through the software development life cycle. Continuous delivery best practices include: • Make every change releasable • Embrace trunk-based development • Deliver through an automated pipeline • No downtime • Release at the granularity of the text

Course 1: Building Cloud Native Applications

Key Enabling Technologies for a Cloud Native Solution (Continued)

Authentication	Developers often need to authenticate users and decide if they can access the system and what level of access they will have. App ID is a service provided by IBM Cloud that allows users to create and use Single Sign On (SSO) solutions for their own applications.
Infrastructure as Code (IaC)	IaC is using code to manage and provision infrastructure. Infrastructure includes things like networks, virtual machines, load balancers, clusters, services, and network connections. Using IaC allows developers to automate the provisioning process using scripts. Benefits of IaC: • **Improves consistency**: New environments or infrastructure are provisioned reliably. • **Increases speed and decreases cost**: IaC enables a quick infrastructure setup for every environment. This can lead to lower costs as the time to deploy, manage, and maintain environments decreases. • **Enables tracking**: Changes to existing infrastructure are made in code, and the changes are tracked. All changes made to a configuration are traceable. • **Corrects drift**: If a part of the infrastructure is modified manually outside of the code, it can be brought back in line with the desired state on the next execution. IaC Tools Schematics is an IBM Cloud service that delivers IaC tools as a service. Schematics provides the ability to consistently deploy and manage cloud infrastructure environments.

Question 1.

Cloud native applications can be defined as:

A. Applications that can be lifted and shifted to the cloud and run on similar hardware

B. Applications where all necessary code is run on a single instance hosted in the cloud

C. Applications that are designed to run exclusively in the public or private cloud

D. Applications that are available to deploy directly from the IBM Cloud Catalog

Answer C. Cloud native applications can be defined as applications that are designed to run exclusively in the public or private cloud.

Question 2.

Independently deployable is a characteristic of a microservices architecture. What is another?

A. Single common database schema

B. Ease of application scaling

C. Complex application dependencies

D. Limited available tools

Answer B. Independently deployable and ease of application scaling are two characteristics of a microservices architecture.

Question 3.

The goal of the twelve-factor app methodology for creating cloud native applications is to:

A. Use a declarative approach to programming.

B. Use an imperative approach to programming.

C. Minimize the need to scale up quickly by deploying replicas.

D. Minimize portability between execution environments.

Answer A. Use a declarative approach to programming is the goal of the twelve-factor app methodology for creating cloud native applications.

Question 4.

A cloud computing execution model in which servers are transparent to the end user is one way to describe serverless architecture. Which of these statements also applies to serverless architecture?

A. Scaling is the responsibility of the user, while provisioning, patching, and maintenance are the responsibilities of the cloud provider.

B. Servers are not used when application code is executed.

C. Available resources are fixed to ensure a consistent pricing model.

D. Resources are automatically scaled up or down based on demand, requiring no user intervention.

Answer D. A cloud computing execution model in which servers are transparent to the end user, and resources are automatically scaled up or down based on demand, requiring no user intervention are two ways to describe serverless architecture.

Question 5.

Application programming interfaces (APIs) enforce a set of rules that allow applications or devices to communicate with each other is one way to describe an API. What is another way to describe it?

A. APIs perform rolling updates and compile code hosted on a developer's local machine.

B. APIs are a means of deploying application updates through automated pipelines.

C. APIs construct HTML code for front-end web applications in a microservices architecture.

D. APIs can be used to communicate data and commands between microservices.

Answer D. Application programming interfaces (APIs) enforce a set of rules that allow applications or devices to communicate with each other, and APIs can be used to communicate data and commands between microservices are two ways to describe an API.

Question 6.

Messaging and event streaming services perform which role in a microservices architecture?

A. Deploy code automatically into each phase of the development cycle.

B. Broadcast state changes to other services such as error logging services.

C. Deploy and manage the availability of microservices across a cluster of host servers.

D. Create and use Single Sign On (SSO) solutions for hosted applications.

Answer B. Broadcast state changes to other services such as error logging services is the role that messaging and event streaming services perform in a microservices architecture.

Question 7.

A client needs to implement a continuous delivery pipeline as a component of their DevSecOps processes. They are evaluating a cloud native approach to deploying several new microservices. What could they consider to be a best practice of continuous delivery on IBM Cloud?

A. Deliver code to a virtual server instance similar to the developer's local machine.

B. Prepare software for production in a single, but holistic development cycle.

C. Deliver code to various environments through automated pipelines.

D. Use multiple branches to avoid version control conflicts when possible.

Answer C. Beliver code to various environments through automated pipelines is considered a best practice of continuous delivery on IBM Cloud.

Course 2: Modernizing Existing Applications

Introduction and Objectives

In Course 2 of the study guide, the subject matter:

- Explores methods used to modernize and rebuild existing applications, enabling them to run in a cloud environment.

Lessons

- Application Modernization Approaches

- Modernize or Rebuild Existing Applications

- Summary

- Knowledge Check

Objective

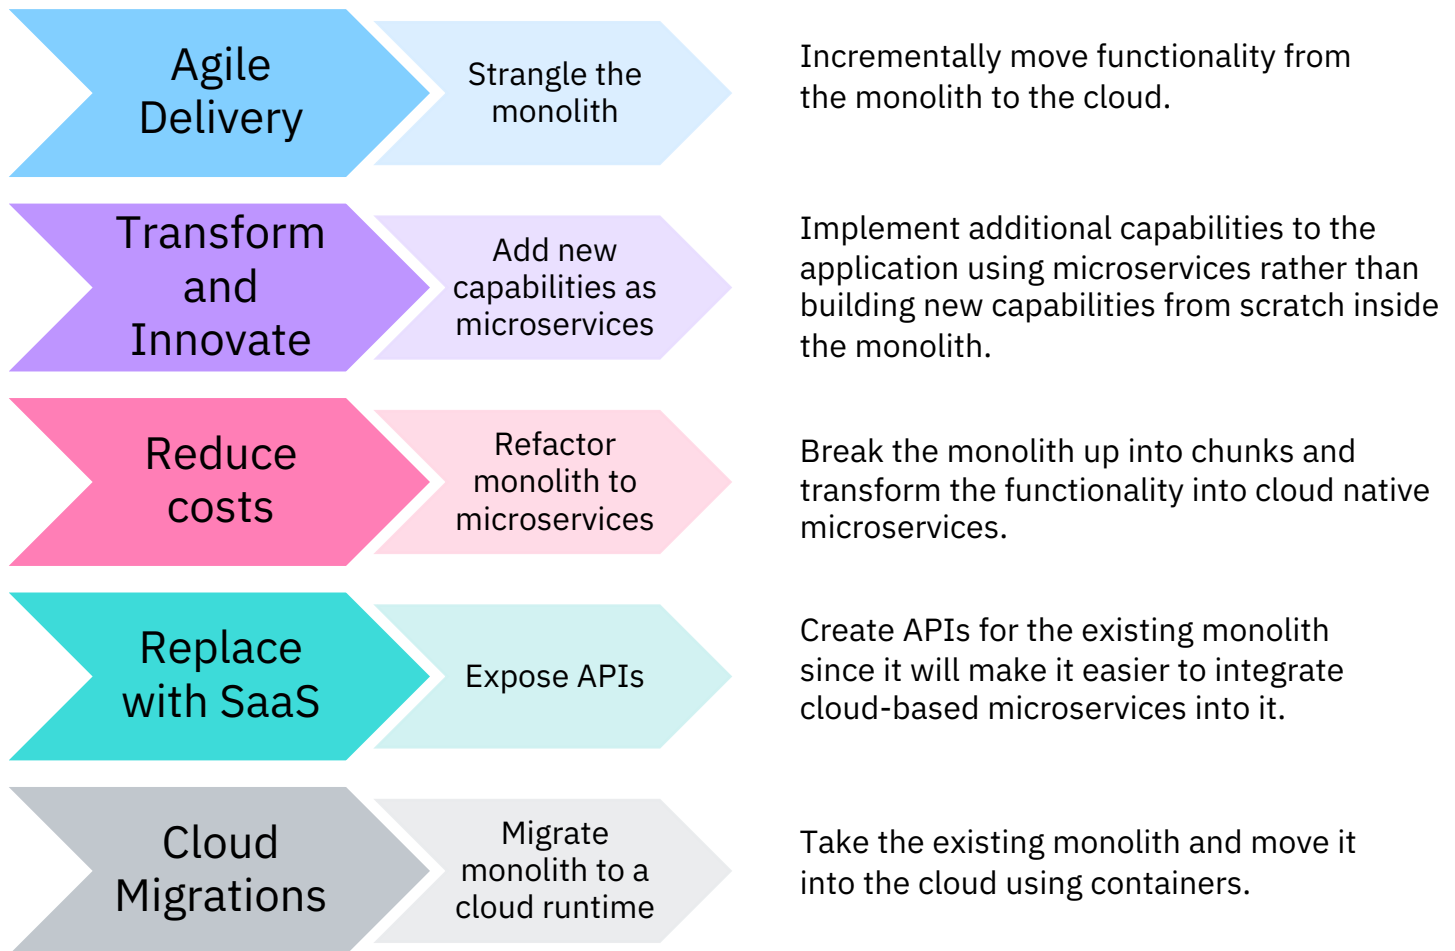- Describe the benefits of modernizing existing applications

Course 2: Modernizing Existing Applications

Modernization Journey
There are a variety of ways to move monolithic applications to the cloud. Each of these patterns can be used to move applications into the cloud with a minimal amount of disruption to their operations.

Modernization Goals **Modernization Patterns**

Modernization Goals	Modernization Patterns	
Agile Delivery	Strangle the monolith	Incrementally move functionality from the monolith to the cloud.
Transform and Innovate	Add new capabilities as microservices	Implement additional capabilities to the application using microservices rather than building new capabilities from scratch inside the monolith.
Reduce costs	Refactor monolith to microservices	Break the monolith up into chunks and transform the functionality into cloud native microservices.
Replace with SaaS	Expose APIs	Create APIs for the existing monolith since it will make it easier to integrate cloud-based microservices into it.
Cloud Migrations	Migrate monolith to a cloud runtime	Take the existing monolith and move it into the cloud using containers.

Course 2: Modernizing Existing Applications

Modernization

Depending on the nature of the existing monolithic application, modernizing an existing application by moving functionality to cloud-based microservices may be the best step in moving to the cloud. By modernizing an application instead of rebuilding it for cloud, development costs can be saved while still enjoying the benefits of a cloud-based application. This approach is best when some of the functionality of the existing application can be moved to microservices.
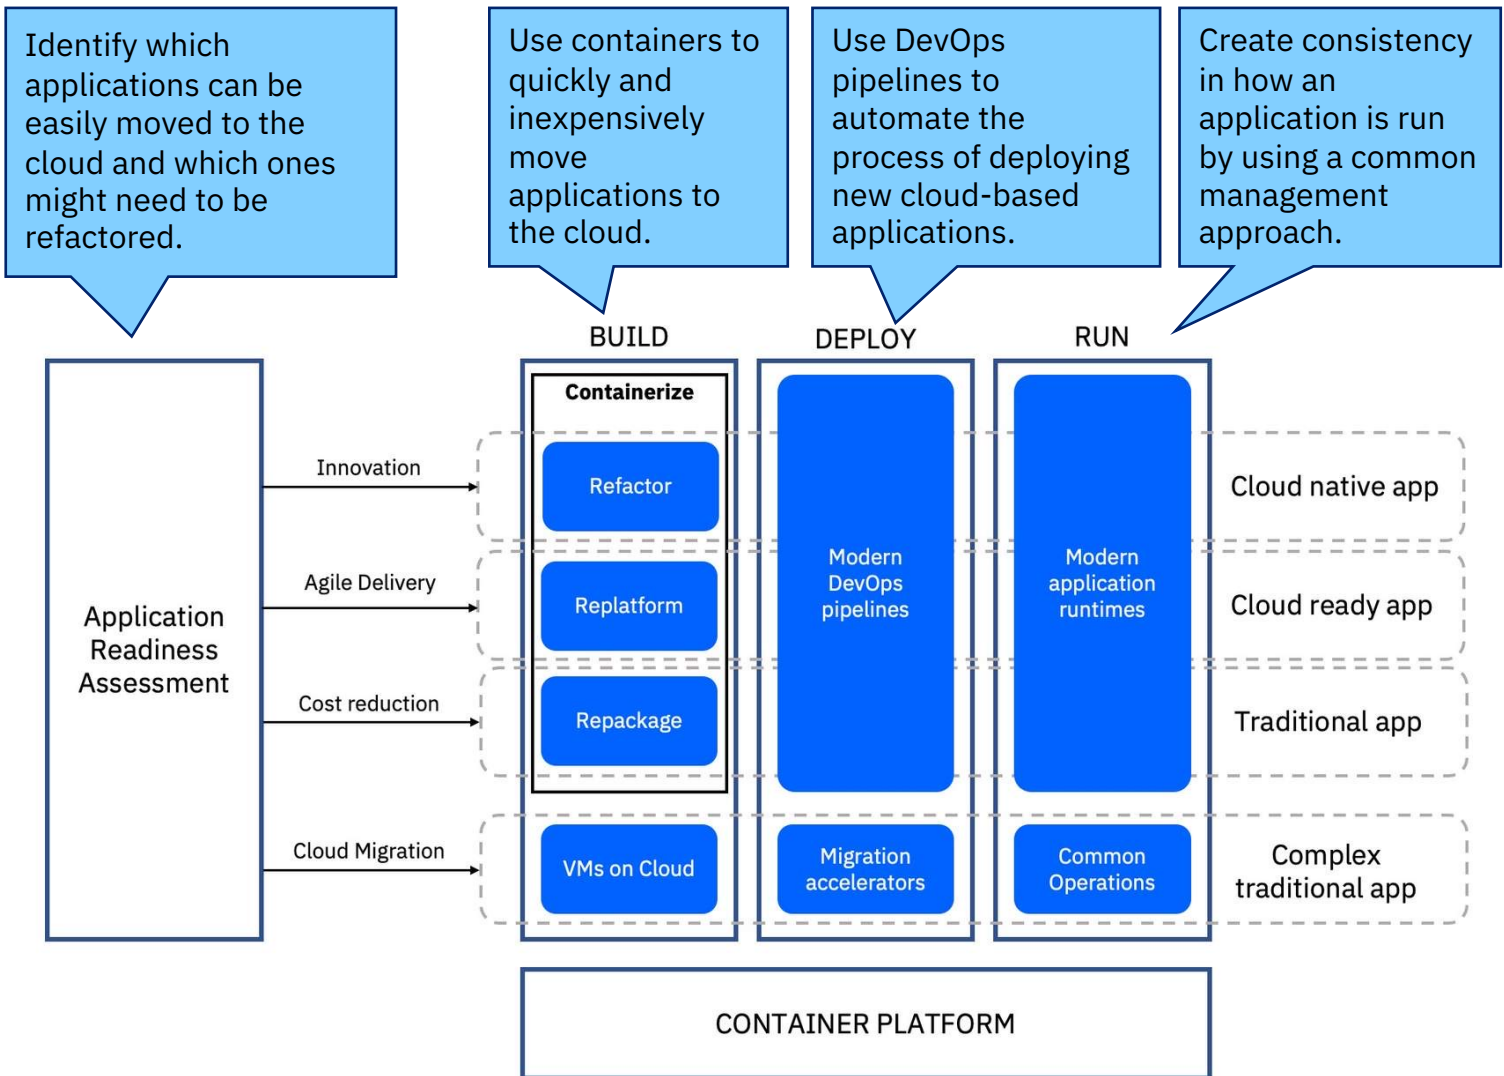
In some cases, it may be necessary to completely rebuild an existing application from scratch. It could be that the existing application needs many updates, or its functionality cannot be easily converted to microservices. In this case, it is necessary to completely rebuild the application from scratch. The process of rebuilding an existing application need not be time consuming and costly since cloud-native applications consist of microservices that have already been developed.

Question 1.

How is an application modernization pattern described?

A.   Break the monolith up into chunks and transform the functionality into cloud native microservices.

B.   Provision new cloud-based virtual servers and storage and migrate the monolith as-is to the cloud.

C.   Combine multiple applications into a single code set and deploy using continuous delivery practices.

D.   Build each new application using microservices on a Kubernetes platform with DevOps.

Answer A. An application modernization pattern is described as breaking the monolith up into chunks and transforming the functionality into cloud native microservices.

Question 2.

What describes application refactoring?

A. When new functionality is built separately in the cloud using a microservices approach and connects using APIs

B. When the underlying server infrastructure is virtualized, and the application is migrated as-is to the cloud, enhancing RTO

C. When a new application is built using cloud native services to completely replace the existing one, which is then retired

D. When an existing monolithic application is broken into chunks, and the functionality is transformed into cloud native microservices

Answer D. Application refactoring is when an existing monolithic application is broken into chunks, and the functionality is transformed into cloud native microservices.

Acronyms

Acronym	Acronym Expansion
API	Application Programming Interface
DevSecOps	Development, Security, and Operations
IaC	Infrastructure as Code
IDE	Integrated Development Environment
IdP	Identity Provider
OS	Operating System
REST	Representational State Transfer
SSO	Single Sign On
SQL	Structured Query Language