

IBM Cloud Advocate v2

IBM C1000-142

Version Demo

Total Demo Questions: 27

Total Premium Questions: 271

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

Topic Break Down

Topic	No. of Questions
Topic 1, Cloud Computing Fundamentals	97
Topic 2, IBM Cloud Essentials	18
Topic 3, IBM Cloud Account Essentials	11
Topic 4, IBM Cloud Security	32
Topic 5, IBM Cloud Services	113
Total	271

QUESTION NO: 1

Write three characteristic of microservices architecture? (Select Three)

- A. Code can be updated more easily
- B. Teams can use different stacks and different programming languages for different components
- C. All components and functions are in a single instance
- D. Components can be scaled independently of one another

ANSWER: A B D

Explanation:

Microservices architecture structures an application as a collection of small services that are loosely coupled and can be developed, deployed, and operated independently. This separation means that code can be updated more easily: a team can modify and release a specific service without requiring a coordinated release of the whole application. It also enables teams to use different stacks and different programming languages for different components when those choices best suit each service's requirements, provided the services communicate through well-defined interfaces or APIs.

Components can be scaled independently of one another because each service can be deployed and allocated resources separately. For example, a service experiencing unusually high request volume can be scaled without scaling unrelated application capabilities. This supports more efficient use of infrastructure and allows teams to align capacity with the workload of each individual service. These characteristics arise from service autonomy, independent deployment, and API-based interaction—the core design principles of microservices. IBM describes microservices as independently deployable services organized around business capabilities and communicating through lightweight mechanisms. See [IBM: What are microservices?](#) and [IBM Cloud documentation: Microservices](#).

QUESTION NO: 2

What is an advantage an enterprise would realize from using VPC?

- A. IBM manages virtual machine operating systems in VPC
- B. Logical isolation from other public cloud tenants, creating a private, secure place on the public cloud
- C. VPC's automatically connect with VPC's hosted by other cloud providers
- D. No egress charges for data leaving IBM Cloud VPC

ANSWER: B

Explanation:

Logical isolation from other public cloud tenants, creating a private, secure place on the public cloud is a core advantage of IBM Cloud Virtual Private Cloud (VPC). A VPC provides a logically isolated virtual network within IBM Cloud, so an enterprise can deploy workloads in a dedicated network environment rather than directly sharing network space with unrelated tenants. Within that environment, the enterprise controls essential network constructs such as subnets, IP address ranges, routing, network access control lists, and security groups. This enables organizations to segment applications, restrict traffic according to their security requirements, and establish private connectivity patterns for workloads handling sensitive data. IBM Cloud VPC is designed to combine the flexibility and on-demand characteristics of public cloud infrastructure with the network isolation and control that enterprises expect for production systems. The isolation is logical rather than a separate physical cloud, but it provides a secure boundary for the organization's cloud resources and supports a defense-in-depth approach when paired with identity and access management, security groups, and appropriate routing rules. IBM describes VPC as a logically isolated virtual network for deploying cloud resources; see the [IBM Cloud VPC overview](#) and the [IBM Cloud Virtual Private Cloud product page](#).

QUESTION NO: 3

Which of the following practices help improve the resilience of an application deployed in IBM Cloud? (Select Three)

- A. Deploying across multiple availability zones
- B. Using load balancing
- C. Backing up persistent data
- D. Running all production components on one virtual server

ANSWER: A B C

Explanation:

Deploying across multiple availability zones is correct because distributing workload components across zones can reduce the effect of an availability-zone failure. A resilient design avoids placing all critical application components in one failure domain.

Using load balancing is correct because a load balancer can distribute requests across healthy application instances. This supports high availability and enables applications to continue serving traffic if an individual instance becomes unavailable.

Backing up persistent data is also correct because backups support recovery from data loss, corruption, accidental deletion, and other events that affect stored information.

Resilience requires a combination of architecture, operational processes, monitoring, and recovery planning. See the [IBM: What is high availability?](#) and the [IBM Cloud high availability and disaster recovery guidance](#).

QUESTION NO: 4

Which of the following are cloud computing characteristics that can help organizations control costs? (Select Three)

- A. On-demand self-service
- B. Rapid elasticity
- C. Measured service
- D. Fixed capacity that cannot be changed after provisioning

ANSWER: A B C

Explanation:

On-demand self-service can help control costs because resources can be provisioned when they are needed instead of being purchased far in advance. **Rapid elasticity** enables capacity to scale with workload demand, reducing the need to maintain infrastructure sized permanently for peak usage. **Measured service** supports consumption-based billing by monitoring resource use and associating charges with consumed capacity.

Cloud computing does not guarantee that every workload has a fixed monthly cost, and it does not remove the need for organizations to monitor usage and apply cost-management practices. See [IBM: What is cloud computing?](#) for an overview of cloud characteristics.

QUESTION NO: 5

Kubernetes relies on ____ configuration files to maintain the application state? (Fill in the Blank)

- A. YAML

ANSWER: A

Explanation:

YAML is correct because Kubernetes uses declarative manifests to describe the desired state of application resources, such as Pods, Deployments, Services, ConfigMaps, and volumes. A YAML manifest expresses resource metadata, specifications, and configuration in a structured, human-readable format. When the manifest is submitted to the Kubernetes API, Kubernetes stores the desired state and continuously works to reconcile the running environment with that declaration. For example, a Deployment YAML file can specify the container image, replica count, labels, and update strategy; Kubernetes then maintains the requested number of replicas as conditions change. This declarative approach supports repeatable deployments, version control, automated delivery pipelines, and consistent configuration across environments. Although Kubernetes resource definitions can also be represented in JSON, YAML is the standard and most commonly used format in Kubernetes documentation, command-line workflows, and operational practice. IBM Cloud Kubernetes Service follows this same Kubernetes manifest model when users deploy and manage workloads on clusters. Kubernetes explains that objects are typically created from YAML manifest files, and IBM documents deployment workflows that apply YAML configuration files to IBM Cloud Kubernetes Service clusters. See the [Kubernetes objects documentation](#) and [IBM Cloud Kubernetes Service application deployment documentation](#).

QUESTION NO: 6

Select the answer that correctly completes the sentence. Blockchain and AI have a ____ relationship with the cloud.

- A. Symbiotic
- B. Internal
- C. Dual
- D. Dependent

ANSWER: A

Explanation:

Symbiotic correctly describes the relationship because cloud computing, artificial intelligence, and blockchain can reinforce one another's capabilities. Cloud platforms provide the elastic compute capacity, storage, networking, and managed services needed to develop, train, deploy, and scale AI workloads without requiring organizations to build all underlying infrastructure themselves. The cloud can also make blockchain networks and applications more accessible by supplying deployable infrastructure and services for building, operating, and integrating distributed-ledger solutions.

In the other direction, AI can derive insights from data associated with cloud-hosted applications and can help automate operational tasks. Blockchain contributes a shared, tamper-evident record that can increase trust, traceability, and transparency among participants in multiparty business processes. When these technologies are used together on cloud infrastructure, each addresses capabilities that complement the others: cloud supplies scalable delivery, AI supplies intelligence, and blockchain supplies trusted shared records. IBM describes blockchain as a way to establish trust among multiple parties and provides blockchain offerings that can be deployed through cloud-based environments. See IBM's [Blockchain overview](#) and [Artificial intelligence overview](#).

QUESTION NO: 7

Which of following statements are true for IBM Cloud® Function? (Select Two)

- A. You have to pay only for the exact time for which your action runs
- B. IBM Cloud® Function able to be scale automatically
- C. IBM Cloud® Function Is a self-contained unit that can run anywhere
- D. Physical infrastructure resources like virtual server, network, storage disks are not required for IBM Cloud® Function

ANSWER: A B

Explanation:

“You have to pay only for the exact time for which your action runs” is correct because IBM Cloud Functions uses a serverless, consumption-based model. A function action is invoked in response to an event or request, and billing is based on the resources consumed while that action executes, rather than on maintaining continuously provisioned servers. This lets teams run code without paying for idle server capacity.

“IBM Cloud® Function able to be scale automatically” is also correct. IBM Cloud Functions is built on Apache OpenWhisk and provides automatic scaling for action invocations. As demand increases, the platform can run additional action instances to process concurrent work; when demand subsides, the platform removes the need for the customer to manage that execution capacity. This operational model is a core benefit of serverless functions: developers focus on deploying actions and connecting triggers, while the cloud platform handles the underlying execution environment and scaling behavior. IBM's Cloud Functions documentation describes actions as code invoked by events and explains the service's serverless execution model. See [IBM Cloud Functions documentation](#) and [IBM Cloud Functions actions documentation](#).

QUESTION NO: 8

True Or False. IBM Watson Studio allows you to train, deploy, and manage your AI models, and prepare and analyze information during a single integrated environment?

- A. TRUE
- B. FALSE

ANSWER: A

Explanation:

TRUE is correct. IBM Watson Studio is an integrated environment for the AI and machine-learning lifecycle. It provides tools for preparing and analyzing data, building and training models, and operationalizing models for deployment and ongoing management. Teams can use its collaborative projects and available tooling to work through data-science tasks in one platform rather than moving between disconnected services for every stage of development. Watson Studio supports common model-development approaches, including visual tools, notebooks, and automated machine learning, helping users move from raw data to trained models efficiently. When a model is ready for use, IBM watsonx.ai/Watson Studio capabilities integrate with deployment and governance-oriented services in the IBM watsonx platform so models can be put into production and monitored as part of their lifecycle. This makes the statement an accurate high-level description of Watson Studio's purpose: providing a unified workspace for data preparation, analysis, model development, training, and operationalization. IBM describes Watson Studio as a platform for building, running, and managing AI models, with collaborative tools for the data-science workflow. See the [IBM Watson Studio product page](#) and the [IBM Cloud Pak for Data as a Service documentation overview](#).

QUESTION NO: 9

When using the Cost Estimator tool, what three formats can be used to download the completed estimate?

- A. CSV
- B. DOC
- C. PDF
- D. PPT
- E. XLSX

ANSWER: A C E

Explanation:

The completed estimate can be downloaded in CSV, PDF, and XLSX formats. These choices support the common needs that arise after configuring estimated IBM Cloud service costs. CSV provides a plain, structured comma-separated file that can be imported into many financial, reporting, and data-analysis tools. XLSX provides an Excel workbook format, making it

useful when the estimate needs further spreadsheet-based review, calculations, formatting, or sharing with finance stakeholders. PDF provides a fixed-layout document that is appropriate for distributing or retaining a readable version of the estimate without requiring recipients to manipulate the underlying data. Together, these export formats enable both machine-readable analysis and business-ready presentation of projected costs. IBM Cloud describes the Cost Estimator as a tool for creating and managing estimates for planned IBM Cloud resources and explains estimate export/download capabilities in its account and billing documentation. See the [IBM Cloud Cost Estimator documentation](#) and the [IBM Cloud billing documentation](#) for current product guidance.

QUESTION NO: 10

What service is the OpenShift platform built on?

- A. Tekton
- B. Jenkins
- C. Kubernetes
- D. TravisCI

ANSWER: C

Explanation:

Kubernetes is correct because Red Hat OpenShift is an enterprise Kubernetes application platform built on the upstream Kubernetes project. Kubernetes provides the foundational container-orchestration capabilities: deploying and scheduling containerized workloads, maintaining their desired state, scaling applications, service discovery, and managing cluster resources. OpenShift incorporates Kubernetes and adds an integrated set of enterprise platform capabilities around it, such as developer workflows, image management, networking, observability, security controls, and automated installation and lifecycle management. Consequently, when IBM Cloud offers Red Hat OpenShift clusters, the cluster's core orchestration layer remains Kubernetes, with OpenShift supplying the broader managed application-platform experience. This relationship is central to understanding OpenShift: applications are packaged in containers and run under Kubernetes control, while OpenShift provides opinionated tools and operational features that help teams build, deploy, secure, and operate those applications consistently. Red Hat describes OpenShift as a Kubernetes platform, and IBM Cloud documentation identifies Red Hat OpenShift on IBM Cloud as a managed OpenShift service for deploying and managing containerized workloads. See [Red Hat OpenShift overview](#) and [IBM Cloud: Red Hat OpenShift overview](#).

QUESTION NO: 11

What are the key tenets of a hybrid cloud?

- A. Interoperability, scalability, and portability
- B. On-demand self-service, controlled zone network access, and measured service
- C. Usable, collaborative, and extensible
- D. Logging, auditing, and network isolation
- E. Interoperability, scalability, portability, and security

ANSWER: E

Explanation:

The key tenets of a hybrid cloud are interoperability, scalability, portability, and security. A hybrid cloud connects and coordinates public-cloud and private-cloud resources so that applications, data, and processes can operate across environments according to business, technical, regulatory, and performance requirements. Interoperability is essential because the environments must be able to communicate and work together through compatible technologies, APIs, integration, and management practices. Portability enables applications and workloads to be moved or deployed across

cloud and on-premises environments with reduced rework and vendor dependency. Scalability allows organizations to expand or reduce resources as demand changes, including using public-cloud capacity when appropriate. Security is equally fundamental: a hybrid strategy must consistently protect data, identities, workloads, and network communications across every connected environment while supporting governance and compliance requirements. IBM describes hybrid cloud as an approach that combines public and private cloud environments, enabling organizations to use each environment for the workloads to which it is best suited. IBM also emphasizes the value of flexibility, integration, and secure operations across these environments. See [IBM: What is hybrid cloud?](#) and [IBM Hybrid Cloud](#).

QUESTION NO: 12

Which IBM Cloud support plan allows opening a case with Security 3 ranking?

- A. Platinum
- B. Subscription
- C. Premium
- D. Unlimited

ANSWER: C

Explanation:

Premium is correct because it is IBM Cloud's highest standard support-plan tier and includes the case-management entitlements needed for security-related support requests, including a case categorized with a Security 3 ranking. IBM Cloud uses rankings to communicate the impact and urgency of a support issue so that cases can be routed and handled according to the applicable support process. A Security 3 case represents a security-support classification that requires the Premium support entitlement when opening the case through IBM Cloud Support. Premium also provides the broadest level of technical support coverage and faster target response objectives for qualifying cases, making it the appropriate plan for customers that require direct security-case access in addition to production technical support. IBM documents the available support plans, their features, and their case eligibility in its support-plan documentation. The process for creating and managing cases, including selecting the relevant issue classification, is described in the IBM Cloud Support Center documentation. See [IBM Cloud support plans](#) and [Getting help and opening IBM Cloud support cases](#).

QUESTION NO: 13

True Or False. Content Delivery Network is a distributed server network that delivers temporarily stored, or cached, copies of website content to users based on the users' geographical location?

- A. True
- B. False

ANSWER: A

Explanation:

True is correct. A content delivery network (CDN) is a geographically distributed network of servers that caches and delivers web content from locations closer to end users. When a user requests a website asset such as an image, video, stylesheet, script, or downloadable file, the CDN can serve a cached copy from a nearby edge server rather than requiring every request to travel back to the origin server. Proximity reduces network latency and can improve page-load speed, responsiveness, availability, and the ability to handle traffic at scale. CDN caching is especially valuable for static or cacheable content that is requested repeatedly by users in different regions. IBM Cloud describes its CDN service as using a globally distributed network to accelerate delivery of internet content to users, while IBM's cloud-computing guidance similarly identifies CDNs as distributed servers that provide content based on user location. This description therefore accurately captures the core purpose and operation of a CDN. See [IBM: What is a CDN?](#) and [IBM Cloud Internet Services documentation](#).

QUESTION NO: 14

Which of the following are common capabilities of IBM Cloud Object Storage? (Select Three)

- A. Stores unstructured data as objects
- B. Uses buckets to organize objects
- C. Provides access through an API
- D. Attaches directly to a virtual server as a boot volume

ANSWER: A B C

Explanation:

Stores unstructured data as objects is correct because object storage stores data as independent objects that include the data, metadata, and a unique identifier. This model is appropriate for large volumes of content such as images, videos, archives, backups, logs, and data sets.

Uses buckets to organize objects is correct. Buckets are containers used to store and organize objects, and bucket-level configuration can define settings such as location, resilience, access controls, and retention behavior. **Provides access through an API** is also correct because applications interact with IBM Cloud Object Storage through supported APIs, including S3-compatible interfaces.

Object storage is different from block storage because it is not presented to a server as a raw disk volume for mounting as a conventional file system. See the [IBM Cloud Object Storage overview](#) and the [IBM Cloud Object Storage API documentation](#).

QUESTION NO: 15

Which of the following include the common goals of application modernization? Select all that apply.

- A. Microservices architecture
- B. DevOps
- C. Cloud migration
- D. Portability

ANSWER: A B C D

Explanation:

Application modernization commonly seeks to make applications easier to change, deploy, scale, and operate in hybrid-cloud environments. **Microservices architecture** supports this goal by decomposing applications into independently deployable services, allowing teams to update and scale functionality with greater agility. **DevOps** is central to modernization because it connects development and operations practices, promotes automation and continuous delivery, and shortens the time required to move reliable changes into production. **Cloud migration** is also a frequent modernization objective, particularly when organizations move workloads to cloud infrastructure or redesign them to use cloud-native platform services. Finally, **Portability** is a common goal because modernized applications are often designed to run consistently across on-premises, private-cloud, and public-cloud environments. Containerization and orchestration platforms help provide this flexibility while reducing dependence on a single underlying environment. These goals are complementary: cloud migration supplies an appropriate operating environment, microservices and DevOps improve application delivery and maintainability, and portability supports hybrid-cloud choice and operational consistency. IBM describes application modernization as updating existing applications and their infrastructure to leverage cloud-native technologies, and identifies cloud-native practices such as containers, microservices, and DevOps as important elements. See [IBM: Application modernization](#) and [IBM: Cloud-native](#).

QUESTION NO: 16

Which of the following are benefits of using a load balancer for an application? (Select Two)

- A. Distribute incoming traffic across multiple application instances
- B. Improve application availability by routing around unhealthy instances
- C. Replace the need for persistent data storage
- D. Automatically remove all application security vulnerabilities

ANSWER: A B

Explanation:

Distributing incoming traffic across multiple application instances is a primary load-balancing benefit. Instead of sending all requests to one server, a load balancer can direct traffic across healthy backend instances, helping the application use available capacity more efficiently.

Improving application availability is also correct. Health checks can identify unhealthy backend instances, allowing the load balancer to stop routing requests to those instances while continuing to direct traffic to healthy ones. This design supports resilient application architectures when multiple instances are deployed.

A load balancer does not permanently store application data and does not eliminate the need to secure application endpoints. See [IBM Cloud load balancers for VPC documentation](#).

QUESTION NO: 17

Select the answer that correctly completes the sentence. ____ is a utility computing billing method that is implemented in cloud computing and geared toward organizations and end users. customers pay for services as they are procured and/or used.

- A. Pay-As-You-Go
- B. Monthly Billing
- C. Fixed billing
- D. Reserved

ANSWER: A

Explanation:

Pay-As-You-Go correctly describes the cloud utility-computing billing model in which customers are charged according to the resources and services they consume. Rather than committing to a large up-front purchase of infrastructure, an organization can provision cloud services when needed and incur charges based on actual use, such as compute capacity, storage consumption, network traffic, or service requests. This model is geared toward both organizations and individual users because it aligns costs with demand: resources can be scaled up during periods of higher activity and scaled down when they are no longer required. IBM Cloud describes its pricing as consumption-based for many services, with charges determined by the selected service plan and measured usage. This approach supports financial flexibility, enables experimentation without substantial capital expenditure, and helps teams avoid paying for unused capacity. The National Institute of Standards and Technology also identifies measured service as an essential cloud-computing characteristic, noting that cloud resource use can be monitored, controlled, and reported to provide transparency for both the provider and consumer. See [IBM Cloud billing and usage documentation](#) and [NIST SP 800-145](#).

QUESTION NO: 18

What was the earliest implementation of cloud computing?

- A. Providing compute resources for complex problems such as asteroid tracking.
- B. Indexing large datasets across universities
- C. LAN parties and cyber cafes
- D. Selling un-used time slice of computing

ANSWER: D

Explanation:

Selling un-used time slice of computing is correct because cloud computing's roots lie in time-sharing systems. Beginning in the 1960s, organizations with expensive mainframes made otherwise idle processing capacity available to multiple users. Rather than each user or organization buying and operating a separate computer, they could consume a portion of centrally managed compute time as needed. This established the core economic and operational model that later evolved into cloud services: pooled infrastructure, shared resources, metered or allocated consumption, and access to computing capacity without owning the underlying hardware. Modern cloud computing expands this historical approach through virtualization, automation, internet-based access, and rapid self-service provisioning, but the foundational concept of making unused computer time available for consumption came first. IBM's overview of cloud-computing history identifies time-sharing as an important precursor to today's cloud model, while the NIST cloud definition describes the modern model's on-demand access to a shared pool of configurable resources. See [IBM: History of cloud computing](#) and [NIST SP 800-145: The NIST Definition of Cloud Computing](#).

QUESTION NO: 19

Which of the following are types of Document Database? (Select Two)

- A. MongoDB
- B. Db2 Hosted
- C. My SQL
- D. Db2
- E. Cloudant

ANSWER: A E

Explanation:

MongoDB and Cloudant are document database technologies. MongoDB is a document-oriented database that stores data as flexible, JSON-like BSON documents rather than requiring a rigid row-and-column schema. This model lets applications represent nested and evolving data structures naturally, which is useful for workloads such as content management, product catalogs, mobile applications, and operational analytics. MongoDB documentation describes documents as the basic unit of data and collections as groups of documents, reflecting its document data model.

IBM Cloudant is IBM's fully managed, distributed NoSQL database service based on the Apache CouchDB ecosystem. It stores and retrieves JSON documents and provides HTTP-based APIs, replication, and synchronization capabilities. Its document-oriented approach is designed for applications that need flexible schemas, high availability, and distributed data access, including web and mobile scenarios. IBM specifically identifies Cloudant as a JSON document database service, making it an appropriate example of a document database alongside MongoDB. See the [MongoDB documentation on the document data model](#) and the [IBM Cloudant getting-started documentation](#).

QUESTION NO: 20

Which IBM Cloud® Pak provides tools to automate the implementation of business policies?

- A. IBM Cloud® Pak for Business Automation

- B. IBM Cloud® Pak for Data
- C. IBM Cloud® Pak for Integration
- D. IBM Cloud® Pak for Security

ANSWER: A

Explanation:

IBM Cloud® Pak for Business Automation is the correct choice because it is IBM's platform for automating business operations, including the decisions and rules that implement business policies. Its business-automation capabilities bring together workflow automation, decision automation, content services, process mining, and document processing. In particular, decision automation supports the management and execution of business rules so that policy-driven decisions can be applied consistently across business processes. This enables organizations to turn policies into governed, repeatable operational logic, update that logic as policies change, and use it within automated workflows. IBM positions Cloud Pak for Business Automation as a modular, containerized platform that helps organizations design, automate, and manage end-to-end business processes and decisions on Red Hat OpenShift. This focus directly matches the need to automate implementation of business policies rather than merely analyze data, connect applications, or manage security events. See IBM's [Cloud Pak for Business Automation product page](#) and its [IBM documentation for Cloud Pak for Business Automation](#) for product capabilities and deployment information.

QUESTION NO: 21

Which IBM cloud service tracks and measure outcomes from AI throughout its life cycle.

- A. Watson Studio
- B. Watson life cycle manager
- C. Watson OpenScale
- D. Watson Knowledge Catalog

ANSWER: C

Explanation:

Watson OpenScale is the IBM service designed to monitor, evaluate, and govern AI model outcomes throughout the model lifecycle. It provides ongoing visibility into deployed model behavior, helping teams assess whether predictions remain accurate, fair, explainable, and aligned with organizational governance requirements. Its monitoring capabilities can identify model drift, such as when production data or model performance changes over time, and can evaluate fairness by detecting potential bias in outcomes for monitored attributes. Watson OpenScale also supports explainability by providing insight into factors that influenced an individual model prediction, enabling users to understand and communicate how an AI-assisted decision was reached. These capabilities are central to responsible AI operations because they allow organizations to move beyond initial model development and continuously measure model quality and trustworthiness after deployment. The service is part of IBM's watsonx governance capabilities and can be used to operationalize AI governance, risk monitoring, and lifecycle oversight across supported model environments. IBM describes Watson OpenScale as a component for monitoring and evaluating AI models for fairness, quality, and explainability. See the [IBM Watson OpenScale product page](#) and the [IBM documentation for Watson OpenScale](#).

QUESTION NO: 22

Which of the following are capabilities of Kubernetes? (Select Three)

- A. Automated scheduling of container workloads
- B. Self-healing of failed container instances
- C. Horizontal scaling of applications

D. Automatic conversion of monolithic applications into microservices

ANSWER: A B C

Explanation:

Kubernetes is an open-source container orchestration platform that automates the deployment, scaling, and management of containerized workloads. **Automated scheduling** is a core capability because Kubernetes places containers on suitable nodes according to declared resource requirements and cluster capacity. **Self-healing** is also supported because Kubernetes can restart failed containers, replace unhealthy instances, and maintain the desired number of replicas.

Horizontal scaling enables an application to increase or decrease the number of running instances in response to workload demand.

Kubernetes does not replace the need to write application code or eliminate the need for an operating system on the worker nodes. See the [Kubernetes concepts overview](#) and [IBM: What is Kubernetes?](#) for further details.

QUESTION NO: 23

Which of the following statements are true for service IDs in IBM Cloud IAM? (Select Two)

- A. They are used by applications and services to access IBM Cloud resources
- B. They can be assigned IAM access policies
- C. They are intended for employees to sign in interactively to the console
- D. They automatically have administrator access to every resource in an account

ANSWER: A B

Explanation:

Service IDs are used by applications and services is correct. A service ID represents a non-human identity that can be assigned IAM access policies, allowing an application, automation tool, or workload to access IBM Cloud resources without using an individual user identity.

Service IDs can be assigned access policies is also correct. Policies define the roles and resource scopes available to the service ID, enabling least-privilege access for automated workloads. API keys can be created for a service ID so that an application can authenticate to IBM Cloud APIs.

A service ID is not intended for interactive console access by an employee, and it does not replace an IBM Cloud account owner. See the [IBM Cloud documentation for service IDs](#).

QUESTION NO: 24

Which service connects a virtual private cloud to IBM services on a private network?

- A. Transient endpoints
- B. Service endpoints
- C. Virtual private endpoints
- D. Elastic endpoints

ANSWER: C

Explanation:

Virtual private endpoints is correct because IBM Cloud Virtual Private Endpoints for VPC provide private connectivity from resources in an IBM Cloud VPC to supported IBM Cloud services. Rather than sending traffic over the public internet, a virtual private endpoint creates a private IP address within a VPC subnet. Workloads such as virtual servers can use that address to reach the target IBM service across IBM's private network.

This design is useful when an organization needs to keep service traffic isolated from public network paths, apply VPC security controls, and meet internal security or compliance expectations. Endpoint access can be controlled through security groups and access management policies, while private DNS resolution can allow applications to use the service's normal hostname and resolve it to the endpoint's private address. IBM Cloud documents virtual private endpoints as the private-network access mechanism for supported services, with each endpoint associated with a VPC, subnet, and target service. See the [IBM Cloud documentation on Virtual Private Endpoints for VPC](#) and the [IBM Cloud guide to provisioning a virtual private endpoint](#).

QUESTION NO: 25

Which applications can be used on any cloud platform and rely on microservices and continuous integration and continuous delivery (CI/CD)?

- A. Cloud Enabled
- B. Cloud Ready
- C. Cloud Base
- D. Cloud Native

ANSWER: D

Explanation:

Cloud Native is correct because cloud-native applications are designed and built specifically to take advantage of cloud computing principles, rather than simply being moved from traditional infrastructure to a cloud environment. They commonly use a microservices architecture, in which independently deployable services can be developed, scaled, and updated separately. This architectural approach supports portability across cloud platforms when applications are packaged and orchestrated with cloud-native technologies such as containers and Kubernetes.

CI/CD is also a foundational cloud-native practice. Continuous integration automates the frequent building and testing of code changes, while continuous delivery or deployment automates the release process. Together, these practices let teams deliver application changes rapidly and reliably. IBM describes cloud-native development as an approach that uses technologies such as microservices, containers, Kubernetes, and DevOps practices to build scalable applications for cloud environments. The Cloud Native Computing Foundation likewise identifies microservices and automation as central aspects of cloud-native systems. See [IBM: What is cloud native?](#) and [Cloud Native Computing Foundation definition](#).

QUESTION NO: 26

Which of the following are characteristics of an access group in IBM Cloud IAM? (Select Two)

- A. It groups users and service identities with similar access requirements
- B. Policies assigned to the access group apply to its members
- C. It provides a private network boundary for resources
- D. It automatically creates a separate IBM Cloud account

ANSWER: A B

Explanation:

It groups users and service identities with similar access requirements is correct. Access groups simplify administration by allowing an organization to define access once for a group rather than creating and maintaining equivalent policies for each individual identity.

Policies assigned to the access group apply to its members is also correct. When a user or service identity is added to an access group, it receives the access granted through that group's policies. This makes access groups useful for job-function-based access management and helps support consistent application of least-privilege permissions.

An access group does not replace a resource group. Resource groups organize cloud resources, whereas access groups organize identities that need similar access. See the [IBM Cloud access groups documentation](#).

QUESTION NO: 27

Which statements are true for Hypervisor? (Select Max)

- A.** A small software layer that enables multiple operating systems to run alongside each other.
(Correct)
- B.** Hypervisor separates the VMs logically, assigning each its own slice of underlying computing power, memory, and storage.
- C.** A Hypervisor is an exact copy of a regular computer or server.
- D.** Hypervisor is a service for virtualization used in OS to deliver softwares in containers with libraries, and configuration files.

ANSWER: A B

Explanation:

A hypervisor is the virtualization layer that makes it possible for several separate operating systems to run on the same physical host. It abstracts the host hardware and presents virtual hardware to each virtual machine, enabling each guest operating system to operate in an isolated environment. This capability is fundamental to server virtualization because it allows organizations to consolidate workloads that would otherwise require separate physical machines.

“Hypervisor separates the VMs logically, assigning each its own slice of underlying computing power, memory, and storage” is also correct. The hypervisor allocates and manages physical resources, including processor capacity, memory, storage, and networking, for individual virtual machines. Logical isolation helps ensure that each virtual machine has its own configured resources and operating environment while sharing the same underlying host infrastructure. IBM describes virtualization as using software to create virtual versions of computing resources and identifies the hypervisor as the software layer that enables multiple virtual machines on one physical system. See IBM's [hypervisor overview](#) and its [virtualization overview](#) for further details.