DUMPSBOSS.

IBM Cloud Professional Architect v6

IBM C1000-172

Version Demo

Total Demo Questions: 7

Total Premium Questions: 74

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Architecture and Design	19
Topic 2, IBM Cloud Services	37
Topic 3, Security	17
Topic 4, Migration	1
Total	74

QUESTION NO: 1

Which two are the private registry options available to IBM Cloud OpenShift clusters?

A. GitHub

B. IBM Cloud Container Registry

C. Docker Hub

D. The internal registry setup at cluster creation time

E. Red Hat Quay

ANSWER: B D

Explanation:

IBM Cloud Container Registry is a private, managed registry service that is designed for storing and distributing container images within IBM Cloud. It integrates with IBM Cloud Identity and Access Management, supports private namespaces, and allows Red Hat OpenShift on IBM Cloud workloads to pull images using the appropriate registry credentials and image-pull secrets. This makes it suitable for centrally managing images that are built in IBM Cloud CI/CD workflows or deployed across IBM Cloud clusters.

The internal registry setup at cluster creation time is also a private-registry choice for a Red Hat OpenShift on IBM Cloud cluster. OpenShift provides an integrated image registry that can store images inside the cluster environment when it is configured with suitable persistent storage. Because this registry is associated with the cluster, it supports OpenShift-native image streams and build/deployment workflows without requiring images to be hosted in a separate external registry service. IBM documents the registry choices and configuration considerations for OpenShift clusters, while IBM Cloud Container Registry documentation describes private namespaces, access control, and image usage. See IBM Cloud Docs: Setting up a private registry and IBM Cloud Docs: Container Registry getting started.

QUESTION NO: 2

What is the name of the program that IBM Cloud follows to ensure its services meet the security and compliance standards of the US government?

A. CIS

B. NIST

C. FedRAMP

D. FIPS

ANSWER: C

Explanation:

FedRAMP is correct because the Federal Risk and Authorization Management Program is the U.S. government-wide program for assessing, authorizing, and continuously monitoring cloud products and services used by federal agencies. It establishes a standardized risk-management process based on NIST security requirements, enabling agencies to evaluate whether a cloud service offering has the controls needed to protect federal information at an appropriate impact level. IBM Cloud maintains FedRAMP-authorized offerings intended to support U.S. federal workloads, and its compliance documentation identifies FedRAMP as a supported government compliance program. FedRAMP authorization is not merely a statement of general security alignment: it requires an independent security assessment, review by the relevant authorizing body, and ongoing continuous monitoring after authorization. This standardized authorization approach helps government agencies reuse security assessments and make informed decisions when adopting cloud services. The FedRAMP Program Management Office describes the program and its authorization process at FedRAMP.gov. IBM's public compliance information, including details on its FedRAMP-authorized cloud environment and available documentation, is available through the IBM Cloud compliance center.

QUESTION NO: 3

An application running only on private VPC subnets must access a supported IBM Cloud service without using public IP addresses or sending traffic through the public internet.

Which two design elements are needed for this private service-access pattern?

Select two.

A. A VPC endpoint gateway for the supported service

B. DNS resolution to the private endpoint addresses

C. A public gateway attached to the VPC

D. A Transit Gateway attachment to the VPC

ANSWER: A B

Explanation:

A VPC endpoint gateway for the supported service provides private endpoint addresses in the VPC for accessing that service. **DNS resolution that directs the service hostname to the private endpoint addresses** ensures that application requests use the endpoint gateway rather than a public service endpoint.

A public gateway is used to provide outbound public-internet access for VPC resources and is not required for private service access through a VPE. A Transit Gateway connects attached networks, such as VPCs and classic infrastructure, but does not replace an endpoint gateway for access to a supported IBM Cloud service.

QUESTION NO: 4

Which IBM Cloud database service supports both relational and non-relational data querying?

A. Databases for Redis

B. Databases for Db2

C. Databases for PostgreSQL

D. Databases for etcd

ANSWER: B

Explanation:

Databases for Db2 is correct because Db2 supports traditional relational data management through tables, rows, columns, relationships, and SQL, while also providing native JSON capabilities for semi-structured document data. JSON documents can be stored in Db2 columns and queried using SQL/JSON functions and predicates, so an application can work with relational and non-relational data in the same database platform. This hybrid capability is useful when a workload needs the consistency, governance, and transactional strengths of a relational database but also must accommodate flexible or evolving document structures, such as application metadata, event payloads, or customer attributes. IBM documents Db2 JSON support as enabling JSON data to be stored and queried alongside relational data, rather than requiring a separate document database for those use cases. The IBM Cloud service delivers this Db2 capability as a managed offering, reducing operational effort for provisioning, maintenance, backups, and availability. See IBM's Db2 JSON data documentation and the IBM Cloud Databases for Db2 service page.

QUESTION NO: 5

Which two search capabilities are available in IBM Cloud Log Analysis?

A. Simple search

B. SQL search

C. Wildcard search

D. Field search

E. Multi search

ANSWER: A C

Explanation:

IBM Cloud Log Analysis supports **Simple search** for quickly locating log entries that contain a specified word, value, or text string. This is useful during initial troubleshooting, when an architect or operator needs to find messages such as an error identifier, service name, host name, request value, or application-generated text without constructing a complex query. The search is applied to the log data displayed in the Log Analysis interface, helping users narrow a large log stream to relevant events.

Wildcard search provides flexible pattern matching when the complete search term is unknown or when several related values must be found together. For example, a wildcard can help locate values that share a known prefix or suffix, such as similarly named services, files, request paths, or error-message variants. This capability is particularly valuable for exploratory investigation because real-world log values often include generated identifiers, versioned names, and variable portions that cannot be anticipated exactly.

These capabilities enable efficient interactive log investigation and complement the service's filtering and analysis workflow. IBM Cloud Log Analysis documentation is available through the IBM Cloud Log Analysis documentation; IBM's broader IBM Cloud documentation portal provides service and operational guidance.

QUESTION NO: 6

An architect is tasked with setting up IBM Cloud Object Storage for data with unpredictable usage patterns. Which storage class should the architect select for this use case?

A. Standard

B. Smart Tier

C. Vault

D. Cold Vault

ANSWER: B

Explanation:

Smart Tier is the appropriate IBM Cloud Object Storage class for data whose access pattern is unpredictable or changes over time. Smart Tier is intended for workloads where it is difficult to determine in advance whether objects will be accessed frequently or infrequently. It provides a single storage class that automatically accounts for changing access behavior, allowing the organization to avoid needing to manually classify objects into a fixed access-frequency tier as usage evolves. This makes it especially useful for datasets with fluctuating demand, such as content repositories, analytics data, or application assets with uncertain retrieval patterns.

IBM Cloud Object Storage storage classes are selected based on expected data-access characteristics. Smart Tier specifically addresses unknown or variable access frequency while retaining immediate object availability when needed. Its pricing model is designed to accommodate changing usage patterns, helping balance cost and access requirements without operational effort to move data between classes. IBM identifies Smart Tier as the class for data with unpredictable access patterns in its storage-class documentation. See IBM Cloud Object Storage documentation: storage classes and IBM Cloud Object Storage overview.

QUESTION NO: 7

An application runs on virtual server instances in two subnets in the same IBM Cloud VPC. The security team requires a control that allows outbound HTTPS traffic from the application tier and permits only the corresponding return traffic, without separately defining inbound ephemeral ports.

Which VPC network control should be used?

A. A security group

B. A network ACL

C. A route table

D. A flow log collector

ANSWER: A

Explanation:

A security group is the correct choice because security groups are stateful virtual firewall controls that are associated with VPC network interfaces. When an outbound HTTPS rule permits a connection, response traffic for that established connection is allowed automatically. This avoids the need to define separate inbound rules for ephemeral response ports.

Network ACLs are stateless and apply at the subnet level, so rules must explicitly allow both directions of a traffic flow. Route tables determine where traffic is sent, not whether a connection is permitted, and flow logs provide traffic metadata for analysis rather than enforcing policy.