

Oracle

1Z0-1072-25 Exam

Oracle Cloud Infrastructure 2025 Architect Associate

**Questions & Answers
Demo**

Version: 4.1

Question: 1

A financial firm is designing an application architecture for its online trading platform that should have high availability and fault tolerance. What should the architect do to avoid any costly service disruptions and ensure data durability?

- A. Create a new Object Storage bucket in another region and configure recycle policy to move data every 5 days.
- B. Create a replication policy to send data to a different bucket in another OCI region.
- C. Create a lifecycle policy to regularly send data from the Standard to Archive storage.
- D. Copy the Object Storage bucket to a block volume.

Answer: B

Explanation:

For an online trading platform requiring high availability and fault tolerance, it's critical to ensure data durability and avoid any costly service disruptions. In Oracle Cloud Infrastructure (OCI), Object Storage is often used to store critical data, such as transaction logs or user data, due to its scalability, durability, and reliability.

Option B is the most suitable approach for ensuring data durability and availability across regions.

Here's why:

Cross-Region Replication (CRR): OCI offers a feature called Cross-Region Replication for Object Storage. This feature allows you to automatically and asynchronously replicate objects in a bucket from one OCI region to another. This setup ensures that even if one region experiences a failure, the data is still available in another region, thereby meeting the requirements for high availability and fault tolerance.

Data Durability: By replicating data to another region, you protect against regional outages. OCI guarantees 99.995% availability for replicated data, which is critical for a financial firm's trading platform where data consistency and durability are paramount.

Disaster Recovery: With data replicated in another region, the trading platform can quickly switch to using the data in the secondary region in case of a disaster in the primary region. This setup significantly reduces recovery time objectives (RTO) and ensures business continuity.

Reference:

Oracle Cloud Infrastructure Documentation: Cross-Region Replication for Object Storage

Oracle Whitepaper: High Availability and Disaster Recovery in Oracle Cloud Infrastructure

Explanation of Incorrect Options:

Option A: Creating a new Object Storage bucket in another region and configuring a recycle policy to move data every 5 days does not provide real-time data availability or the fault tolerance required for a financial application. Recycle policies are intended for managing the lifecycle of data, not for high availability or disaster recovery.

Option C: While lifecycle policies are useful for moving less frequently accessed data to a more cost-effective storage tier (e.g., from Standard to Archive), they do not address cross-region redundancy or real-time availability, which are critical for this use case.

Option D: Copying an Object Storage bucket to a block volume is not a recommended practice for ensuring data durability and fault tolerance. Block volumes are used for persistent storage attached to compute instances, and copying object storage data to block volumes does not achieve the same level of redundancy and cross-region availability as replication policies.

Thus, Option B is the correct and most efficient method for ensuring high availability and fault tolerance in this scenario.

Question: 2

Which TWO statements are TRUE about Private IP addresses in Oracle Cloud Infrastructure (OCI)?

- A. By default, the primary VNIC of an instance in a subnet has one primary private IP address and one secondary private IP address.
- B. Each VNIC can only have one private IP address.
- C. By default, the primary VNIC of an instance in a subnet has one primary private IP address.
- D. A private IP can have an optional public IP assigned to it if it resides in a public subnet.

Answer: C, D

Explanation:

In Oracle Cloud Infrastructure (OCI), understanding how private IP addresses work is crucial for configuring network interfaces and managing instances within your Virtual Cloud Network (VCN).

Primary VNIC and Private IP Address:

When an instance is launched in OCI, it is attached to a Virtual Network Interface Card (VNIC). The primary VNIC, which is automatically created during the instance launch, is associated with a primary private IP address by default. This private IP address is essential for the instance to communicate within the VCN. The primary private IP address is automatically assigned and cannot be removed from the primary VNIC while the instance is running. This supports the statement C.

Additional Private IPs:

Contrary to statement B, each VNIC can indeed have multiple private IP addresses, but by default, the primary VNIC comes with only one primary private IP. You can manually add secondary private IPs if needed. However, the additional IPs are not assigned by default; hence, A is incorrect.

Public IP Association:

For instances requiring internet access, a public IP address can be optionally assigned to the private IP address if the instance is in a public subnet. This is critical for scenarios where an instance needs to communicate with the internet or external networks. This aligns with statement D.

Relevant OCI Documentation:

[Oracle Cloud Infrastructure Networking Overview](#)

[VNICs and Private IPs](#)

These references provide additional context and detail on how private IP addresses work within OCI and clarify the correct statements.

Question: 3

What happens to the performance level of a volume when it is detached from an instance?

- A. The performance level is adjusted to Balanced.
- B. The performance level remains unchanged.
- C. The performance level is adjusted to Higher Performance.
- D. The performance level is adjusted to Lower Cost (0 VPUs/GB).

Answer: D

Explanation:

In Oracle Cloud Infrastructure (OCI), when a block volume is detached from an instance, its performance level is automatically adjusted to the "Lower Cost" tier, which provides 0 VPUs (Volume Performance Units) per GB. This adjustment helps reduce costs when the block volume is not actively being used by a compute instance.

Key Points:

Volume Performance Levels: OCI offers various performance tiers for block volumes, including "Higher Performance," "Balanced," and "Lower Cost." These tiers determine the level of IOPS (Input/Output Operations Per Second) and throughput available to the volume.

Automatic Adjustment: When a block volume is detached from an instance, OCI automatically optimizes the cost by switching the volume to the "Lower Cost" performance tier. This tier offers minimal performance, suitable for data that is not actively accessed.

Cost Management: This automatic adjustment is beneficial for managing costs, as it prevents users from incurring unnecessary charges for higher performance levels when the volume is not in use.

Reference:

Oracle Cloud Infrastructure Documentation: Block Volume Performance Levels

Question: 4

Which compute capacity type would you select to meet these requirements?

- A. Capacity reservation
- B. Preemptible capacity
- C. On-demand capacity
- D. Dedicated host

Answer: C

Explanation:

On-demand capacity is the compute capacity type in Oracle Cloud Infrastructure (OCI) that allows

you to provision and use compute instances whenever needed, without any long-term commitment. This flexibility is ideal for various workloads, including development, testing, and production environments, where immediate availability and scalability are crucial.

Key Points:

On-Demand Capacity: On-demand compute instances provide users with the flexibility to spin up instances as required and only pay for the time the instances are running. This model is most suitable for workloads with unpredictable usage patterns or short-term requirements.

Flexibility and Scalability: With on-demand capacity, you can quickly scale your resources up or down based on your application's needs, ensuring that you only pay for the resources you actually use.

No Commitment: Unlike reserved capacity, on-demand capacity does not require any long-term commitment or upfront payment, making it an attractive option for organizations looking to avoid capital expenditures.

Reference:

Oracle Cloud Infrastructure Documentation: OCI Compute Pricing

Oracle Cloud Infrastructure Documentation: Compute Instance Lifecycle

Explanation of Incorrect Options:

A . Capacity reservation: This option allows you to reserve capacity in advance, ensuring that resources are available when needed. It's ideal for predictable workloads but may not be as cost-effective for fluctuating demands.

B . Preemptible capacity: Preemptible instances are a lower-cost option where instances can be terminated by OCI if resources are needed elsewhere. This is suitable for non-critical workloads that can tolerate interruptions.

D . Dedicated host: Dedicated hosts provide physical servers for your exclusive use, offering isolation and predictable performance. This option is more suitable for workloads requiring dedicated resources or compliance needs.

Thus, Option C: On-demand capacity is the correct choice for most general-purpose workloads needing flexibility and immediate availability without long-term commitment.

Question: 5

Which TWO options will accomplish a fully redundant connection from an on-premises data center to a Virtual Cloud Network (VCN) in the us-ashburn-1 region?

- A. Configure a Site-to-Site VPN from a single on-premises CPE.
- B. Configure two FastConnect virtual circuits to the us-ashburn-1 region and terminate them in diverse hardware on-premises.
- C. Configure one FastConnect virtual circuit to the us-ashburn-1 region and the second FastConnect virtual circuit to the us-phoenix-1 region.
- D. Configure one FastConnect virtual circuit to the us-ashburn-1 region and a Site-to-Site VPN to the us-ashburn-1 region.

Answer: B, D

Explanation:

For a fully redundant connection from an on-premises data center to a VCN in the OCI us-ashburn-1

region, it is important to ensure high availability and fault tolerance. Here's how each option contributes to redundancy:

Option B: Two FastConnect Virtual Circuits:

FastConnect provides a dedicated, private connection with higher bandwidth and more consistent performance compared to a VPN. To achieve redundancy, you can configure two FastConnect circuits in the same region (us-ashburn-1), each terminated on diverse hardware on-premises. This setup ensures that even if one circuit or its associated hardware fails, the other circuit can maintain the connection. This ensures no single point of failure in the connectivity to OCI. Thus, option B is correct.

Option D: FastConnect and Site-to-Site VPN:

Another approach to redundancy is to have a mix of connection types. By setting up one FastConnect circuit and one Site-to-Site VPN, both terminating in the same region (us-ashburn-1), you create a diverse connection path. If the FastConnect connection fails, traffic can automatically route through the VPN connection, maintaining connectivity. This setup adds an extra layer of redundancy, making option D correct as well.

Incorrect Options:

Option A: Only configuring a Site-to-Site VPN from a single on-premises CPE does not provide redundancy because it involves just one connection path. If that connection or the CPE fails, there would be no fallback.

Option C: Configuring FastConnect circuits to different regions (us-ashburn-1 and us-phoenix-1) does not provide redundancy within a single region but rather across regions, which is not required for regional redundancy.

Relevant OCI Documentation:

[OCI FastConnect Overview](#)

[Using Site-to-Site VPN and FastConnect for Redundancy](#)

These references offer more detailed information on setting up redundant connections and the benefits of each connection type within OCI.