CERTSWARRIOR

Oracle

1Z0-1151-25

Oracle Cloud Infrastructure 2025 Multicloud Architect Professional

Questions&AnswersPDF

Latest Version: 6.0

Question: 1

When setting up Oracle Interconnect for Google Cloud, what is the primary function of the Cloud Router in Google Cloud?

A. To act as a firewall for traffic between OCI and Google Cloud.
B. To manage IPsec tunnels between OCI and Google Cloud.
C. To exchange routing information with the OCI Dynamic Routing Gateway (DRG) using BGP.
D. To provide network address translation (NAT) for instances in the Google Cloud VPC.

Answer: C

Explanation:
The Cloud Router in Google Cloud is a fully distributed and managed Google Cloud service that provides BGP routing. In the context of Oracle Interconnect for Google Cloud, its primary function is to establish BGP sessions with the OCI Dynamic Routing Gateway (DRG) to exchange routing information. This allows traffic to be dynamically routed between the two cloud environments.
Here's why the other options are incorrect:
A. To act as a firewall for traffic between OCI and Google Cloud. While firewalls are important for security, the Cloud Router's primary function is routing, not firewalling. Firewall rules would be configured separately in both OCI and Google Cloud.
B. To manage IPsec tunnels between OCI and Google Cloud. Oracle Interconnect for Google Cloud, when using Partner Interconnect or Dedicated Interconnect, relies on a direct connection or a partner's network, not IPsec tunnels over the public internet.
D. To provide network address translation (NAT) for instances in the Google Cloud VPC. While Cloud NAT is a separate Google Cloud service that provides NAT, it's not the primary function of the Cloud Router in the context of the Interconnect.

Question: 2

Which of the following is NOT a primary benefit typically associated with adopting a multicloud strategy?

A. Reduced risk of vendor lock-in, fostering greater negotiating power.
B. Access to a wide array of specialized services from different providers, driving innovation.
C. Simplified operational management through a unified, multi-cloud platform.
D. Improved resilience and business continuity by distributing workloads across multiple environments

Answer: C

Explanation:
Here's why:

Multicloud complexity: While there are tools emerging to help with multi-cloud management, the reality is that managing multiple cloud environments increases operational complexity. Each cloud provider has its own set of tools, APIs, management consoles, and best practices. Integrating these disparate systems and achieving true unified management is a significant challenge, not a simplification.
Here's why the other options are typical benefits:
a) Reduced risk of vendor lock-in: By using multiple cloud providers, organizations avoid being completely dependent on a single vendor. This provides greater flexibility, negotiating power, and the ability to switch providers if necessary.
b) Access to a wide array of specialized services: Different cloud providers excel in different areas. A multicloud strategy allows organizations to leverage the best-of-breed services from each provider, optimizing for specific workloads and driving innovation.
d) Improved resilience and business continuity: Distributing workloads across multiple cloud environments provides redundancy and fault tolerance. If one cloud provider experiences an outage, applications can continue to run in other environments, enhancing business continuity.

Question: 3

Which type of traffic is NOT supported by the cross-cloud connection between Oracle Cloud Infrastructure (OCI) and Microsoft Azure?

A. Traffic from an Azure Virtual Network (VNet) to a peered OCI Virtual Cloud Network (VCN) in a different OCI region
B. Traffic between an Azure VNet and OCI VCN using private IP addresses
C. Traffic from an Azure VNet to a peered OCI VCN within the same OCI region
D. Traffic between your on-premises network and the OCI VCN through the Azure VNet

Answer: D

Explanation:
The cross-cloud interconnection between OCI and Azure, facilitated by FastConnect and ExpressRoute, establishes a direct, private connection between the two cloud environments. It does not extend to transitive routing from on-premises networks through Azure to OCI.
Here's a breakdown:
A. Traffic from an Azure Virtual Network (VNet) to a peered OCI Virtual Cloud Network (VCN) in a different OCI region: This is supported. The interconnection allows traffic to flow between VNets and VCNs, even across different regions within each cloud.
B. Traffic between an Azure VNet and OCI VCN using private IP addresses: This is the core functionality of the interconnection. It enables communication using private IP addresses, ensuring secure and private communication between the two clouds.
C. Traffic from an Azure VNet to a peered OCI VCN within the same OCI region: This is also supported. The connection works regardless of whether the VCNs are in the same or different OCI regions.
D. Traffic between your on-premises network and the OCI VCN through the Azure VNet: This is not supported. The interconnection is designed for direct connectivity between Azure and OCI. It doesn't act as a transit point for on-premises traffic to reach OCI. To connect your on-premises network to OCI, you would need a separate connection, such as an OCI FastConnect or a VPN connection directly to OCI. Similarly, to connect your on-premises network to Azure, you would need an Azure ExpressRoute or a VPN connection to Azure

Question: 4

How can an organization ensure secure and efficient data transfer between their frontend data analytics applications in Microsoft Azure and the backend Oracle Autonomous Data Warehouse in Oracle Cloud Infrastructure (OCI) in a multicloud solution?

A. By leveraging a VPN Gateway to create an encrypted tunnel between Azure and OCI for secure data transfer
B. By using public internet connections to transfer data between Azure and OCI, encrypting the data in transit
C. By implementing a hybrid cloud approach that integrates on-premises infrastructure with both Azure and OCI
D. By establishing a dedicated, private connection between Azure and OCI using Azure ExpressRoute and Oracle FastConnect

Answer: D

Explanation:
Here's why:
Azure ExpressRoute and Oracle FastConnect: This combination provides a direct, private connection between the two cloud environments. This offers several advantages:
High bandwidth: Enables fast data transfer, crucial for analytics workloads.
Low latency: Minimizes delays in data transfer, improving application performance.
Enhanced security: Data travels over a private connection, reducing exposure to the public internet.
Reliability: Dedicated connections offer more consistent performance compared to internet-based connections.
Why the other options are less suitable:
a. By leveraging a VPN Gateway to create an encrypted tunnel between Azure and OCI for secure data transfer: While VPNs provide secure connections, they generally offer lower bandwidth and higher latency compared to dedicated connections like ExpressRoute and FastConnect. This can be a bottleneck for data-intensive analytics applications.
b. By using public internet connections to transfer data between Azure and OCI, encrypting the data in transit: While encryption protects the data during transit, using the public internet introduces variability in performance (bandwidth and latency) and potential security risks compared to a private connection. This is not suitable for production analytics workloads that require consistent and reliable performance.
c. By implementing a hybrid cloud approach that integrates on-premises infrastructure with both Azure and OCI: Introducing on-premises infrastructure adds complexity and doesn't directly address the need for efficient data transfer between Azure and OCI. While a hybrid approach might be part of a broader strategy, it's not the primary solution for this specific scenario.

Question: 5

Which scenario is least likely to need Oracle Interconnect for Google Cloud?

A. Migrating data from an on-premises data center to OCI

B. Deploying a multicloud disaster recovery solution
C. Running latency-sensitive applications across OCI and GCP
D. Leveraging OCI’s Autonomous Database with GCP’s AI tools

Answer: A

Explanation:
Here's why:
Oracle Interconnect for Google Cloud (or any cloud interconnect, for that matter) is designed for direct, high-bandwidth, low-latency connections between cloud providers (in this case, OCI and GCP). It's used for scenarios where you need tight integration and communication between workloads running in different clouds.
Migrating data from an on-premises data center to OCI is a one-time data transfer operation. While network connectivity is important for this migration, it doesn't necessitate the ongoing, dedicated connection that Oracle Interconnect provides. You would typically use other methods like:
Data transfer appliances: Physical devices for large-scale data transfer.
Cloud storage services: Uploading data to cloud storage (like OCI Object Storage) and then importing it into other OCI services.
VPN or Direct Connect: Establishing a secure connection between your on-premises network and OCI.
The other options are more likely to require Oracle Interconnect:
Deploying a multicloud disaster recovery solution: If your DR solution involves replicating data or workloads between OCI and GCP, a high-bandwidth, low-latency connection like Oracle Interconnect is crucial for fast failover and recovery.
Running latency-sensitive applications across OCI and GCP: Applications that require real-time communication between components running in different clouds benefit significantly from the low latency provided by Oracle Interconnect.
Leveraging OCI’s Autonomous Database with GCP’s AI tools: If you need to frequently move data between Autonomous Database in OCI and AI/ML services in GCP for processing, Oracle Interconnect would provide the necessary performance.

Question: 6

What is the purpose of a Network Security Group (NSG) in OCI?

A. To define routing policies for subnets within a VCN.
B. To act as a central point for managing internet access for all subnets in a VCN.
C. To provide a virtual firewall at the VNIC level, allowing for granular security policies based on source and destination NSGs.
D. To provide DNS resolution services for instances within a VCN.

Answer: C

Explanation:
Here's a breakdown of Network Security Groups (NSGs) in OCI:

Virtual Firewall at the VNIC Level: NSGs act as virtual firewalls that control traffic at the Virtual Network Interface Card (VNIC) level. This means you can apply security rules to individual instances or groups of instances by associating them with NSGs.

Granular Security Policies: NSGs allow you to define granular security policies based on:

Source and Destination IP addresses or CIDR blocks: You can specify which IP addresses or CIDR blocks are allowed to send or receive traffic.

Source and Destination Ports: You can specify which TCP or UDP ports are allowed.

Protocols: You can specify which protocols (e.g., TCP, UDP, ICMP) are allowed.

Source and Destination NSGs: This is a key feature. You can create rules that allow traffic between specific NSGs, creating micro-segmentation within your VCN.

Why other options are incorrect:

A. To define routing policies for subnets within a VCN: Routing policies are defined by Route Tables, not NSGs.

B. To act as a central point for managing internet access for all subnets in a VCN: Internet access is managed by an Internet Gateway and associated Route Tables, not NSGs. While NSGs can control traffic entering and leaving a subnet via the Internet Gateway, that is not their primary function.

D. To provide DNS resolution services for instances within a VCN: DNS resolution is provided by a DNS Resolver within the VCN, not NSGs.

Question: 7

Suppose you have a highly sensitive dataset stored in an Autonomous Database. What is the BEST way to ensure that no data leaves the EU?

A. Configure a Virtual Cloud Network (VCN) with a security list blocking all outbound traffic.

B. Use Data Masking to obfuscate sensitive data before it leaves the database.

C. Select an EU region for your Autonomous Database deployment and enable customer-managed encryption keys.

D. Implement Transparent Data Encryption (TDE) with a customer-managed key.

Answer: C

Explanation:

Here's why:

EU Region Selection: Deploying the Autonomous Database in an EU region (e.g., Frankfurt, Amsterdam, Paris, Dublin) ensures that the data physically resides within the EU's geographical boundaries. This addresses data residency requirements directly.

Customer-Managed Encryption Keys (CMEK): Enabling CMEK gives you control over the encryption keys used to protect your data at rest. This provides an additional layer of security and helps you meet compliance requirements related to key management. It also ensures that even if someone were to somehow exfiltrate the encrypted data, they wouldn't be able to decrypt it without your keys, which reside in your control, presumably also within the EU.

Why the other options are less suitable:

a) Configure a Virtual Cloud Network (VCN) with a security list blocking all outbound traffic: While blocking all outbound traffic would prevent data from leaving the VCN, it would also make the database practically unusable. Applications and users would not be able to connect to it. This is an overly restrictive and impractical solution.

b) Use Data Masking to obfuscate sensitive data before it leaves the database: Data masking is useful for protecting sensitive data when it needs to be shared with non-production environments or third parties. However, it doesn't prevent data from leaving the EU. The masked data would still be subject to data residency regulations.
d) Implement Transparent Data Encryption (TDE) with a customer-managed key: TDE encrypts data at rest, which is a good security practice. Having a customer-managed key adds further control. However, TDE alone does not guarantee data residency. The data could still be stored in a non-EU region.

Question: 8

You are provisioning Oracle Database@Azure. After creating the Exadata Infrastructure, what is the next essential step within the OCI console related to the database environment?

A. Creating the Azure Virtual Network Peering.
B. Creating the Exadata VM Cluster.
C. Configuring the Azure Network Security Groups.
D. Setting up the cross-connect between OCI and Azure.

Answer: B

Explanation:
After the Exadata Infrastructure (the physical hardware) is provisioned, the next crucial step within the OCI console is creating the Exadata VM Cluster. This VM cluster is the actual compute environment where the Oracle database instances will run. It's a layer of virtualization on top of the physical Exadata hardware. The other options are related to networking and connectivity, which are important but come after the VM cluster is created.

Question: 9

Which of the following scenarios would NOT be a valid example of how multicloud might help in this objective?

A. Leveraging specific cloud providers that offer significant discounts on compute instances during off-peak hours.
B. Using a consistent pricing model for all services regardless of the cloud provider.
C. Deploying workloads on cloud providers that offer better pricing for specific services used by the application.
D. Employing a multi-cloud container orchestration platform to effectively utilize existing resources across different environments.

Answer: B

Explanation:
Here's why:
Pricing variability is key: A core principle of multicloud cost optimization is taking advantage of the differences in pricing models between cloud providers. If all providers had consistent pricing, there

would be no cost advantage to choosing one over another for specific services or usage patterns. The ability to shop around and select the most cost-effective provider for a given workload or time of day is a key driver of multicloud cost savings.
Here's why the other options are valid examples:
a) Leveraging specific cloud providers that offer significant discounts on compute instances during off-peak hours: This is a classic example of multicloud cost optimization. By shifting workloads to providers with lower off-peak pricing, organizations can significantly reduce compute costs.
c) Deploying workloads on cloud providers that offer better pricing for specific services used by the application: Different cloud providers have different pricing strategies for various services. By analyzing these differences and deploying workloads accordingly, organizations can optimize costs. For example, one provider might have cheaper storage, while another has cheaper data transfer.
d) Employing a multi-cloud container orchestration platform to effectively utilize existing resources across different environments: Container orchestration platforms like Kubernetes can help optimize resource utilization across multiple clouds. By efficiently scheduling containers based on resource availability and cost, organizations can minimize waste and reduce overall cloud spending.

Question: 10

To minimize the cost of your Autonomous Database while ensuring it scales automatically to handle peak loads. Which configuration is MOST suitable?

A. Autonomous Database on Dedicated Infrastructure with manual scaling.
B. Autonomous Database on Shared Infrastructure with auto-scaling enabled.
C. Base Database on a virtual machine with auto-scaling configured.
D. Base Database on bare metal with manual scaling.

Answer: B

Explanation:
Here's why:
Shared Infrastructure: Using shared infrastructure is inherently more cost-effective than dedicated infrastructure. You share the underlying hardware resources with other users, which reduces your costs significantly.
Auto-scaling: Enabling auto-scaling allows the Autonomous Database to automatically adjust its compute and storage resources based on workload demands. This means you only pay for the resources you actually consume. During peak loads, the database scales up to handle the increased demand, and during periods of low activity, it scales down to minimize costs.
Why other options are less suitable:
a) Autonomous Database on Dedicated Infrastructure with manual scaling: Dedicated infrastructure is the most expensive option. Manual scaling requires manual intervention to adjust resources, which is not ideal for handling unpredictable peak loads and does not minimize cost automatically.
c) Base Database on a virtual machine with auto-scaling configured: While this offers some flexibility, it requires significantly more management overhead compared to Autonomous Database. You are responsible for managing the operating system, database software, backups, and other administrative tasks. Autonomous Database handles all of this for you.
d) Base Database on bare metal with manual scaling: Bare metal is the most expensive and least flexible option. Manual scaling makes it unsuitable for handling peak loads efficiently and cost-effectively.