

# Oracle

## 1Z0-997-20 Exam

**Oracle Cloud Infrastructure 2020 Architect Professional Exam**

**Questions & Answers  
(Demo Version – Limited Content)**

**Thank you for Downloading 1Z0-997-20 exam PDF Demo**

# Version: 4.0

---

## Question: 1

---

You have deployed a web application targeting a global audience across multiple Oracle Cloud Infrastructure (OCI) regions.

You decide to use Traffic Management Geo-Location based Steering Policy to serve web requests to users from the region closest to the user. Within each region you have deployed a public load balancer with 4 servers in a backend set. During a DR test disable all web servers in one of the regions however, traffic Management does not automatically direct all users to the other region.

Which two are possible causes?

- A. You did not setup a Route Table associated with load Balancer's subnet
- B. You did not setup an HTTP Health Check associated with Load Balancer public IP in the disabled region.
- C. Rather than using Geo-Location based Steering Policy, you should use Failover Policy Type to serve traffic.
- D. One of the two working web servers in the other region did not pass its HTTP health check
- E. You did not correctly setup the Load Balancer HTTP health check policy associated with backend set

---

**Answer: B, E**

---

Explanation:

Managing Traffic Management GEOLOCATION Steering Policies

Geolocation steering policies distribute DNS traffic to different endpoints based on the location of the end user. Customers can define geographic regions composed of originating continent, countries or states/provinces (North America) and define a separate endpoint or set of endpoints for each region.

The Health Checks service allows you to monitor the health of IP addresses and hostnames, as measured from geographic vantage points of your choosing, using HTTP and ping probes. After configuring a health check, you can view the monitor's results. The results include the location from which the host was monitored, the availability of the endpoint, and the date and time the test was performed.

Also you can Combine Managing Traffic Management GEOLOCATION Steering Policies with Oracle Health Checks to fail over from one region to another

The Load Balancing service provides health status indicators that use your health check policies to report on the general health of your load balancers and their components.

if you misconfigure the health check Protocol between the Load balancer and backend set that can lead to not get an accurate response as example below

If you run a TCP-level health check against an HTTP service, you might not get an accurate response. The TCP handshake can succeed and indicate that the service is up even when the HTTP service is ly

configured or having other issues. Although the health check appears good customers might experience transaction failures.

---

### Question: 2

Your team is conducting a root analysis (RCA) following a recent, unplanned outage. One of the block volumes attached to your production WebLogic server was deleted and you have tasked with identifying the source of the action. You search the Audit logs and find several Delete actions that occurred in the previous 24 hours. Given the sample of this event.

```
"event":{
  "tenantId":"ocidl.tenancy.ocl..aaaaaaaamp6954bqkimnbuciasaaaaa"
  "compartmentId":"ocidl.compartment.ori..aaaaaaaav4xtwimindk7znpuAlaaa"
  "compartmentName":"Production"
  "eventId":"14a87512 dbrillej,AD6-041027d191/9"
  "eventName":"DeleteVolume"
  "eventSource":"BlockVolumes"
  "eventType":"ServiceAPI"
  "principalId":"ocidl.user.ocl..aaaaaaaig1SKke1b62pz3ualqwy6otzd7daaqaaaaa"
  "credentialId":""
  "requestAction":"DELETE"
  "requestId":"csid06406dcb4a7999cecId51604ce52/r79253t1#lthilb36blad34bMS1040/PA112B68FfOK3011165F6SUM00"
  "requestAgent":"Mozilla/5." (Windows NT 10.0; *Win64; x64) ApploWebKit/531.36 WM, like Gecko) Chrome/15.0.377.14..."
  "requestHeaders":{...
}
"requestOrigin":"129.254.11.219"
"request Resource":"/20160918/volumes/ocidl.volume.ocl.iad.abuwcljtxksq424tohccipilbzz13w)rrij2ezissSeal05125kzxliq"
"responseStatus":"204"
```

Which item from the event log helps you identify the individual or service that initiated the DeleteVolume API call?

- A. requestAgent
- B. eventSource
- C. principalId
- D. requestOrigin
- E. eventId

---

**Answer: C**

---

Explanation:

The Oracle Cloud Infrastructure Audit service automatically records calls to all supported Oracle Cloud Infrastructure public application programming interface (API) endpoints as log events. Currently, all services support logging by Audit.

Every audit log event includes two main parts:

- Envelopes that act as a container for all event messages
- Payloads that contain data from the resource emitting the event message

The identity object contains the following attributes.

data.identity.authType The type of authentication used.

data.identity.principalId The OCID of the principal.

data.identity.principalName The name of the user or service. This value is the friendly name associated

with principalId .

---

### Question: 3

An online registration system is currently hosted on one large Oracle Cloud Infrastructure (OCI) Bare metal compute Instance with attached block volume to store of the users' data

a. The registration system accepts the information from the user, including documents and photos then performs automated verification and processing to check if the user is eligible for registration.

The registration system becomes unavailable at times when there is a surge of users using the system the existing architecture needs improvement as it takes a long time for the system to complete the processing and the attached block volumes are not large enough to use data being uploaded by the users.

Which is the most effective option to achieve a highly scalable solution?

A. Attach more Block volumes as the data volume increases, use Oracle Notification Service (ONS) to distribute tasks to a pool of compute instances working in parallel, and Auto Scaling to dynamically size the pool of instances depending on the number of notifications received from the Notification Service. Use Resource Manager stacks to replicate your architecture to another region.

B. Change your architecture to use an OCI Object Storage standard tier bucket, replace the single bare metal instance with a Oracle Streaming Service (OSS) to ingest the incoming requests and distribute the tasks to a group of compute instances with Auto Scaling

C. Upgrade your architecture to use a pool of Bare metal servers and configure them to use their local SSDs for faster data access Set up Oracle Streaming Service (OSS) to distribute the tasks to the pool of Bare metal instances with Auto Scaling to dynamically increase or decrease the pool of compute instances depending on the length of the Streaming queue.

D. Upgrade your architecture to use more Block volumes as the data volume increases. Replace the single bare metal instance with a group of compute instances with Auto Scaling to dynamically increase or decrease the compute instance pools depending on the traffic.

---

**Answer: D**

---

---

#### **Question: 4**

---

A data analytics company has been building its next generation big data and analytics platform on Oracle Cloud Infrastructure (OCI). They need a storage service that provides the scale and performance that their big data applications require such as high throughput to compute nodes with low latency file operations in addition, their data needs to be stored redundantly across multiple nodes in a single availability domain and allows concurrent connections from multiple compute instances hosted on multiple availability domains.

Which OCI storage service can you use to meet this requirement?

- A. Object Storage
- B. File System Storage
- C. Archive storage
- D. Block Volume

---

**Answer: B**

---

Explanation:

Oracle Cloud Infrastructure File Storage service provides a durable, scalable, secure, enterprise-grade network file system. You can connect to a File Storage service file system from any bare metal, virtual

machine, or container instance in your Virtual Cloud Network (VCN). You can also access a file system from outside the VCN using Oracle Cloud Infrastructure FastConnect and Internet Protocol security (IPSec) virtual private network (VPN).

Use the File Storage service when your application or workload includes big data and analytics, media processing, or content management, and you require Portable Operating System Interface (POSIX)- compliant file system access semantics and concurrently accessible storage. The File Storage service is designed to meet the needs of applications and users that need an enterprise file system across a wide range of use cases

**Question: 5**

You have an Oracle database system in a virtual cloud network (VCN) that needs to be accessible on port 1521 from your on-premises network CIDR 172.17.0.0/24.

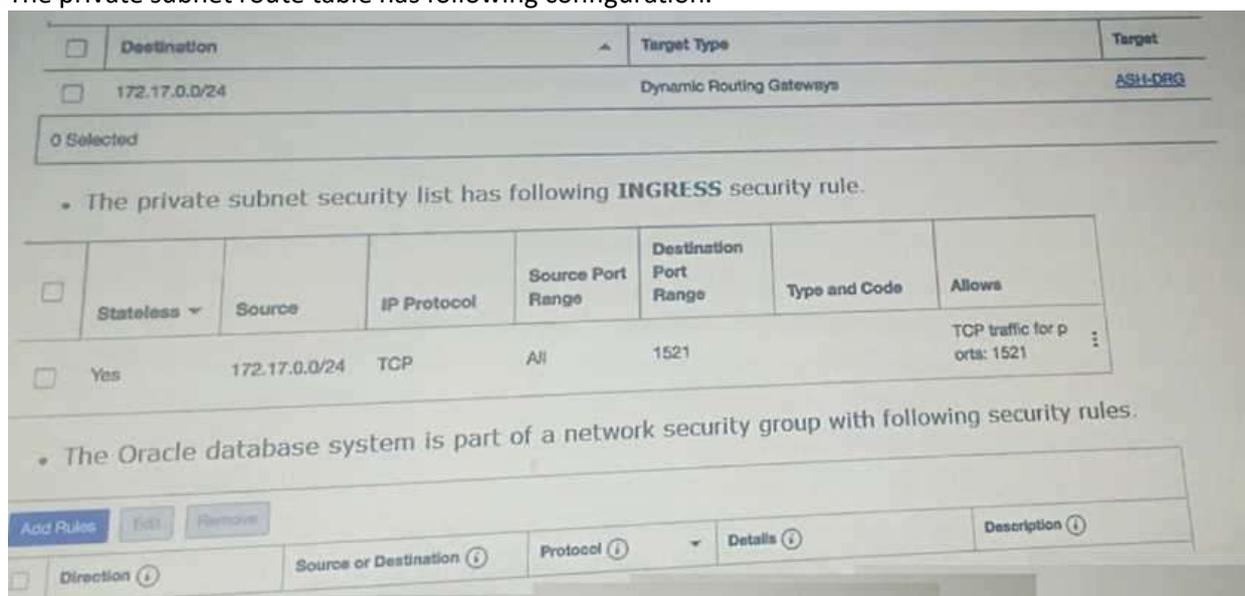
You have the following configuration currently.

Virtual cloud network (VCD) is associated with a Dynamic Routing Gateway (DRG), and DRG has an active IPSec connection with your on-premises data center.

Oracle database system is hosted in a private subnet

The private subnet route table has the following configuration

The private subnet route table has following configuration.



However, you are still unable to connect to the Oracle Database system. Which action will resolve this issue?

- A) Add an EGRESS rule in network security group as following.



- B) Add a route rule in the private subnet route table as following.



C)  
Add an EGRESS rule in private subnet security list as following.

<input type="checkbox"/>	Stateless ▾	Destination	IP Protocol	Source Port Range	Destination Port Range	Type and Code	Allows
<input type="checkbox"/>	Yes	172.17.0.0/24	TCP	1521	All		TCP traffic for port ≤ All

D)  
Add an EGRESS rule in private subnet security list as following.

<input type="checkbox"/>	Stateless ▾	Destination	IP Protocol	Source Port Range	Destination Port Range	Type and Code	Allows
<input type="checkbox"/>	No	172.17.0.0/24	TCP	All	1521		TCP traffic for port ≤: 1521

- A. Option A
- B. Option B
- C. Option C
- D. Option D

**Answer: C**

**Question: 6**

You have provisioned a new VM.DenseIO2.24 compute instance with local NVMe drives. The compute instance is running production application. This is a write heavy application, with a significant Impact to the business if the application goes down. What should you do to help maintain write performance and protect against NVMe devices failure.

- A. NVMe drive have built in capability to recover themselves so no other actions are required
- B. Configure RAID 6 for NVMe devices.
- C. Configure RAID 1 for NVMe devices.
- D. Configure RAID 10 for NVMe devices.

**Answer: D**

Explanation:

VM.DenseIO2.24 compute instance include locally attached NVMe devices. These devices provide extremely low latency, high performance block storage that is ideal for big data, OLTP, and any other workload that can benefit from high-performance block storage.

A protected RAID array is the most recommended way to protect against an NVMe device failure. There are three RAID levels that can be used for the majority of workloads:

RAID 1: An exact copy (or mirror) of a set of data on two or more disks; a classic RAID 1 mirrored pair contains two disks

RAID 10: Stripes data across multiple mirrored pairs. As long as one disk in each mirrored pair is functional, data can be retrieved

RAID 6: Block-level striping with two parity blocks distributed across all member disks If you need the best possible performance and can sacrifice some of your available space, then RAID 10 array is an option.

---

**Question: 7**

---

A civil engineering company is running an online portal In which engineers can upload there constructions photos, videos, and other digital files.

There is a new requirement for you to implement: the online portal must offload the digital content to an Object Storage bucket for a period of 72 hours. After the provided time limit has elapsed, the portal will hold all the digital content locally and wait for the next offload period.

Which option fulfills this requirement?

- A. Create a pre-authenticated URL for the entire Object Storage bucket to read and list the content with an expiration of 72 hours.
- B. Create a pre authenticated URL lot each object that Is uploaded to the Object Storage bucket with an expiration of 72 hours.
- C. Create a Dynamic Group with matching rule for the portal compute Instance and grant access to the Object Storage bucket for 72 hours.
- D. Create a pre authenticated URL for the entire Object Storage bucket to write content with an expiration of 72 hours.

---

**Answer: D**

---

Explanation:

Pre-authenticated requests provide a way to let users access a bucket or an object without having their own credentials, as long as the request creator has permission to access those objects.

For example, you can create a request that lets operations support user upload backups to a bucket without owning API keys. Or, you can create a request that lets a business partner update shared data in a bucket without owning API keys.

When creating a pre-authenticated request, you have the following options:

You can specify the name of a bucket that a pre-authenticated request user has write access to and can upload one or more objects to.

You can specify the name of an object that a pre-authenticated request user can read from, write to, or read from and write to.

Scope and Constraints

Understand the following scope and constraints regarding pre-authenticated requests:

Users can't list bucket contents.

You can create an unlimited number of pre-authenticated requests.

There is no time limit to the expiration date that you can set.

You can't edit a pre-authenticated request. If you want to change user access options in response to changing requirements, you must create a new pre-authenticated request.

The target and actions for a pre-authenticated request are based on the creator's permissions. The request is not, however, bound to the creator's account login credentials. If the creator's login credentials change, a pre-authenticated request is not affected.

You cannot delete a bucket that has a pre-authenticated request associated with that bucket or with an object in that bucket.

---

**Question: 8**

---

You are working as a solution architect for an online retail store to create a portal to allow the users to pay for their groceries using credit cards. Since the application is not fully compliant with the Payment Card Industry Data Security Standard (PCI DSS), your company is looking to use a third party payment service to process credit card payments.

The third party service allows a maximum of 5 public IP addresses at a time. However, your website is using Oracle Cloud Infrastructure (OCI) Instance Pool Auto Scaling policy to create up to 15 instances during peak traffic demand, which are launched in VCN private subnets and attached to an OCI public Load Balancer. Upon user payment, the portal connects to the payment service over the Internet to complete the transaction.

What solution can you implement to make sure that all compute instances can connect to the third party system to process the payments at peak traffic demand?

- A. Route credit card payment request from the compute instances through the NAT Gateway. On the third-party services, whitelist the public IP associated with the NAT Gateway.
- B. Whitelist the Internet Gateway Public IP on the third party service and route all payment requests through the Internet Gateway.
- C. Create an OCI Command Line Interface (CLI) script to automatically reserve public IP address for the compute instances. On the third services, whitelist the Reserved public IP.
- D. Route payment request from the compute instances through the OCI Load Balancer, which will then be routed to the third party service.

---

**Answer: D**

---

Explanation:

You can OCI Load Balancer for this solution which can route the Public IPs of Load balancer to Traffic to third party services which allows a maximum of 5 public IP addresses at a time. However, your website is using Oracle Cloud Infrastructure (OCI) Instance Pool Auto Scaling policy to create up to 15 instances during peak traffic demand.

---

**Question: 9**

---

Your customer recently ordered for a 1-Gbps Fast Connect connection in ap-tokyo-1 region of Oracle Cloud Infrastructure (OCI). They will use this to connect one Virtual Cloud Network (VCN) in their production (OC1) tenancy and VCN in their development OC1 tenancy.

As a Solution Architect, how should you configure and architect the connectivity between on-premises and VCNs in OCI?

- A. Create two private virtual circuits on the FastConnect link. Create two Dynamic Routing Gateways, one for each VCN. Attach the virtual circuits to the dynamic routing gateways.
- B. You cannot achieve connectivity using a single FastConnect link as the production and the development VCNs are in separate tenancies. Request one more FastConnect connection.
- C. Create a single private virtual circuit over FastConnect and attach FastConnect to either of the VCN's Dynamic Routing Gateway. Use Remote Peering to peer production and development VCNs.
- D. Create a hub-VCN that uses Dynamic Routing Gateway (DRG) to communicate with on-premises network over FastConnect. Connect the hub-VCN to the production VCN spoke and with

development VCN spoke, each peered via their respective local Peering Gateway (LPG)

---

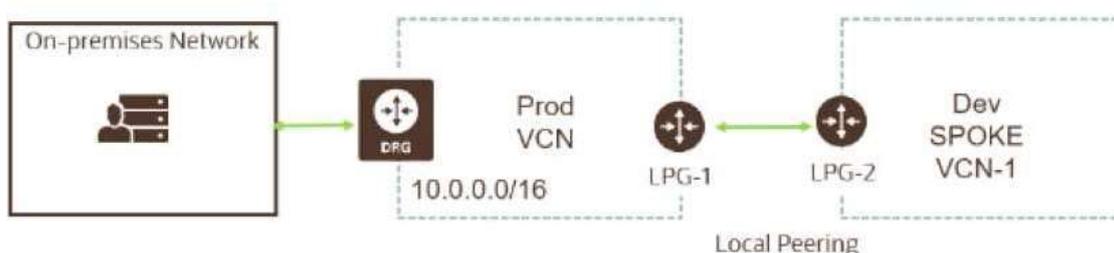
**Answer: D**

---

Explanation:

There's an advanced routing scenario called transit routing that enables communication between an on-premises network and multiple VCNs over a single Oracle Cloud Infrastructure FastConnect or IPSec VPN.

The VCNs must be in the same region and locally peered in a hub-and-spoke layout. As part of the scenario, the VCN that is acting as the hub has a route table associated with each LPG (typically route tables are associated with a VCN's subnets).




---

### Question: 10

---

Which three scenarios are suitable for the Oracle Infrastructure (OCI) Autonomous transaction Processing Server less (ATP-S) deployment?

- A. well established, online auction marketplace is running an application where there is database usage 24x7 but also has peaks of activity that the hard to predict when the peaks happen, the total activities may reach 3 times the normal activity level (Correct)
- B. A small startup is deploying a new application for eCommerce and it requires database to store customers' transactions the team is not sure of what the load will look like since it is a new application. (Correct)
- C. A midsize company is considering migrating its legacy on-premises MongoDB database to Oracle Cloud Infrastructure (OCI). The database has significantly higher workloads on weekends than weekdays
- D. A developer working on an Internal project needs to use a database during work hours but doesn't need it during nights or weekends. the project budget requires her to keep costs low. (Correct)
- E. A manufacturing company is running Oracle E-Business Suite application on-premises. They are looking to move this application to OCI and they want to use a managed database offering for their database tier.

---

**Answer: ABD**

---

Explanation:

MongoDB is a cross-platform document-oriented database program. Classified as a NoSQL database program, MongoDB uses JSON-like documents with schema, so the best to be migrated to Oracle NoSQL Database.

<https://blogs.oracle.com/nosql/migrate-mongodb-data-to-oracle-nosql-database>

Autonomous transaction Processing Serverless (ATP-S) isn't supported yet for EBS database

**Thank You for trying 1Z0-997-20 PDF Demo**

# Start Your 1Z0-997-20 Preparation

**[Limited Time Offer]** Use Coupon “**dumps20**” for extra 20% discount on the purchase of PDF. Test your 1Z0-997-20 preparation with actual exam questions.