



1Z0-997^{Q&As}

Oracle Cloud Infrastructure 2019 Architect Professional

Pass Oracle 1Z0-997 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.passapply.com/1z0-997.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Oracle
Official Exam Center

- ⚙ **Instant Download** After Purchase
- ⚙ **100% Money Back** Guarantee
- ⚙ **365 Days** Free Update
- ⚙ **800,000+** Satisfied Customers





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

Correct Answer: BE

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

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 ls 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.

Correct Answer: D

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 3

A hospital in Austin has hosted its web based medical records portal entirely In Oracle cloud Infrastructure (OCI) using Compute Instances for its web-tier and DB system database for its data tier. To validate compliance with Health Insurance Portability and Accountability (HIPAA), the security professional to check their systems it was found that there are a lot of unauthorized coming requests coming from a set of IP addresses originating from a country in Southeast Asia. Which option can mitigate this type of attack?

- A. Block the attacking IP address by creating by Network Security Group rule to deny access to the compute Instance where the web server ls running



- B. Block the attacking IP address by implementing a OCI Web Application Firewall policy using Access Control Rules
- C. Mitigate the attack by changing the Route table to redirect the unauthorized traffic to a dummy Compute instance
- D. Block the attacking IP address by creating a Security List rule to deny access to the subnet where the web server is running

Correct Answer: B

WAF can protect any internet facing endpoint, providing consistent rule enforcement across a customer's applications. WAF provides you with the ability to create and manage rules for internet threats including Cross-Site Scripting (XSS), SQL Injection and other OWASP-defined vulnerabilities. Unwanted bots can be mitigated while tactically allowing desirable bots to enter. Access rules can limit based on geography or the signature of the request. As a WAF administrator you can define explicit actions for requests that meet various conditions. Conditions use various operations and regular expressions. A rule action can be set to log and allow, detect, or block requests

QUESTION 4

You are working as a solution architect with a global automotive provider who is looking to create a multi-cloud solution. They want to run their application tier in Microsoft Azure while utilizing the Oracle DB Systems in the Oracle Cloud Infrastructure (OCI). What is the most fault tolerant and secure solution for this customer?

- A. Create an Oracle database in OCI Virtual Cloud Network (VCN) and connect to the application tier running in Microsoft Azure over the Internet.
- B. Create a FastConnect virtual circuit and choose Microsoft Azure from the list of providers available to setup Network connectivity between application tier running in Microsoft Azure Virtual Network and Oracle Databases running in OCI Virtual Cloud (VCN)
- C. Use OCI Virtual Cloud Network remote peering connection to create connectivity among application tier running in Microsoft Azure Virtual Network and Oracle Databases running in OCI Virtual Cloud Network (VCN).
- D. Create a VPN connection between the application tier, running in Azure Virtual Network and Oracle Databases running in OCI Virtual Cloud Network (VCN).

Correct Answer: C

Oracle and Microsoft have created a cross-cloud connection between Oracle Cloud Infrastructure and Microsoft Azure in certain regions. This connection lets you set up cross-cloud workloads without the traffic between the clouds going over the internet. You can connect your VNet and VCN so that traffic that uses private IP addresses goes over the cross-cloud connection. For example, the following diagram shows a VNet that is connected to a VCN. Resources in the VNet are running a .NET application that access an Oracle database that runs on Database service resources in the VCN. The traffic between the application and database uses a logical circuit that runs on the cross-cloud connection between Azure and Oracle Cloud Infrastructure. The two virtual networks must belong to the same company and not have overlapping CIDRs. The connection requires you to create an Azure ExpressRoute circuit and an Oracle Cloud Infrastructure FastConnect virtual circuit.

QUESTION 5

An upcoming e-commerce company has deployed their online shopping application on OCI. The application was deployed on compute instances with autoscaling configuration for application servers fronted by a load balancer and OCI Autonomous Transaction Processing (ATP) in the backend. In order to promote their e-commerce platform 50% discount was announced on all the products for a limited period. During the day 1 of promotional period it was observed



that the application is running slow and company's hotline is flooded with complaints. What could be two possible reasons for this situation?

- A. The health check on some of the backend servers has failed and the load balancer has taken those servers temporarily out of rotation
- B. As part of autoscaling, the load balancer shape has dynamically changed to a larger shape to handle more incoming traffic and the system was slow for a short time during this change
- C. The health check on some of the backend servers has failed and the load balancer was rebooting these servers.
- D. The autoscaling has already scaled to the maximum number of instances specified in the configuration and there is no room of scaling

Correct Answer: AD

[Latest 1Z0-997 Dumps](#)

[1Z0-997 Study Guide](#)

[1Z0-997 Braindumps](#)