Oracle

1Z0-997 Exam

Oracle Cloud Infrastructure 2019 Architect Professional Exam

Questions & Answers
Demo

Version: 7.1

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 closets 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

A global retailer is setting up the cloud architecture to be deployed in Oracle Cloud infrastructure (OCI) which will have thousands of users from two major geographical regions: North America and Asia Pacific. The requirements of the services are:
* Service needs to be available 27/7 to avoid any business disruption
* North American customers should be served by application running In North American regions
* Asia Pacific customers should be served by applications running In Asia Pacific regions
* Must be resilient enough to handle the outage of an entire OCI region

A. OCl DNS, Traffic Management with Failover steering policy
B. OCl DNS, Traffic Management with Geolocation steering policy. Health Checks
C. OCl DNS, Traffic Management with Geolocation steering policy
D. OCl DNS,' Traffic Management with Load Balancer steering policy, Health Checks

Answer: B

Explanation:
GEOLOCATION STEERING
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. Combine with Oracle Health Checks to fail over from one region to another

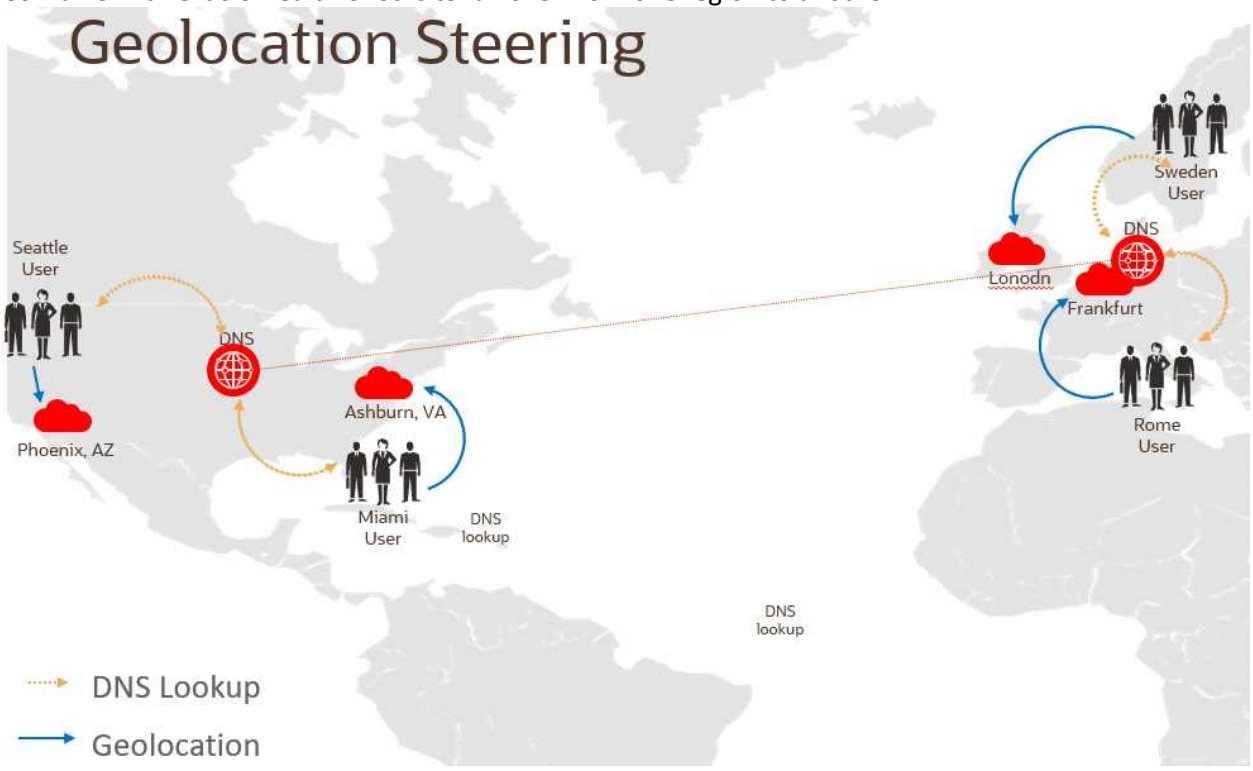

Question: 3

A retail company has several on-premises data centers which span multiple geographical locations. They plan to move some of their applications from on-premises data centers to Oracle Cloud Infrastructure (OCI). For these applications running in OCI, they still need to interact with applications running on their on-premises data centers to Oracle Cloud Infrastructure (OCI). for these applications running in OCI. they still need to interact with applications running on their on-premises data centers. These applications require highly available, fault-tolerant network connections between on premises data centers and OCI.
Which option should you recommend to provide the highest level of redundancy?

A. Oracle cloud Infrastructure provides network redundancy by default so that no other operations are required
B. If your data centers span multiple, geographical locations, use only the specific IP address as a static route for the specific geographical location
C. Set up both IPSec VPN and FastConnect to connect your on premises data centers to Oracle Cloud Infrastructure.
D. Use FastConnect private peering only to ensure secure access from your data center to Oracle Cloud Infrastructure
E. Set up a single IPSec VPN connection (rom your data center to Oracle Cloud Infrastructure since It is cost effective

Answer: B

Explanation:
If your data centers span multiple geographical locations, we recommend using a broad CIDR (0.0.0.0/0) as a static route in addition to the CIDR of the specific geographical location. This broad CIDR provides high availability and flexibility to your network design. For instance, the following diagram shows two networks in separate geographical areas that each connect to Oracle Cloud Infrastructure. Each area has a single on-premises router, so two IPSec VPN connections can be created. Note that each IPSec VPN connection has two static routes: one for the CIDR of the particular geographical area, and a broad 0.0.0.0/0 static route.

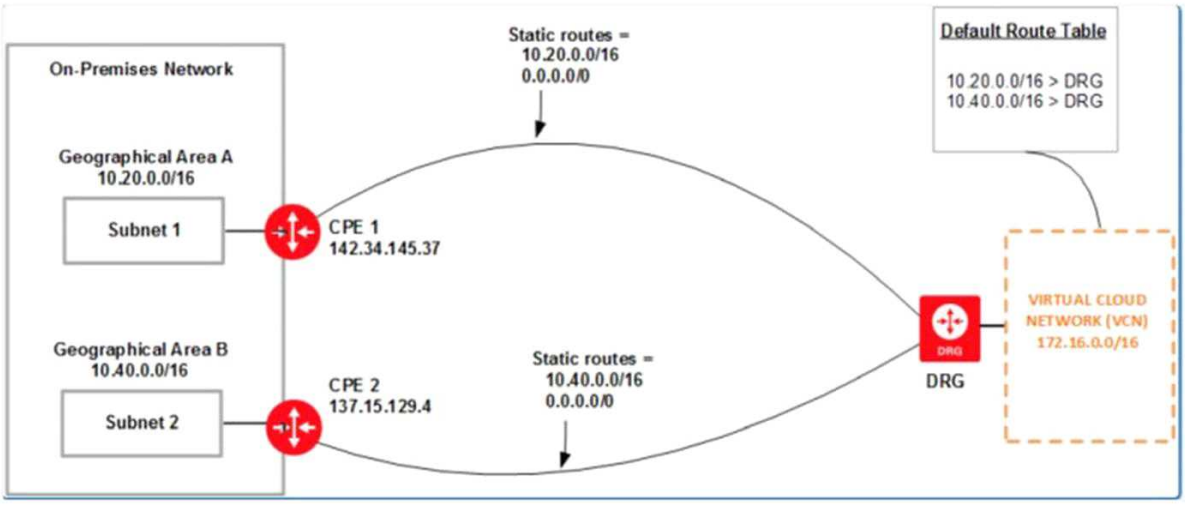

Question: 4

A global retailer has decided to re-design its e-commerce platform to have a micro-services architecture. They would like to decouple application architecture into smaller, independent services using Oracle Cloud Infrastructure (OCI). They have decided to use both containers and servers technologies to run these application instances.
Which option should you recommend to build this new platform?

A. Install a kubernetes cluster on OCI and use OCI event service.
B. Use Oracle Container Engine for kubernetes, OCI Registry and OCI Functions.
C. Use OCI Resource Manager to automate compute Instances provisioning and use OCI Streaming service.
D. Use OCI functions, OCI object storage and OCI event service.

Answer: B

Explanation:
Oracle Functions is a fully managed, multi-tenant, highly scalable, on-demand, Functions-as-a-Service platform. It is built on enterprise-grade Oracle Cloud Infrastructure and powered by the Fn Project open source engine. Use Oracle Functions (sometimes abbreviated to just Functions) when you want to focus on writing code to meet business needs.

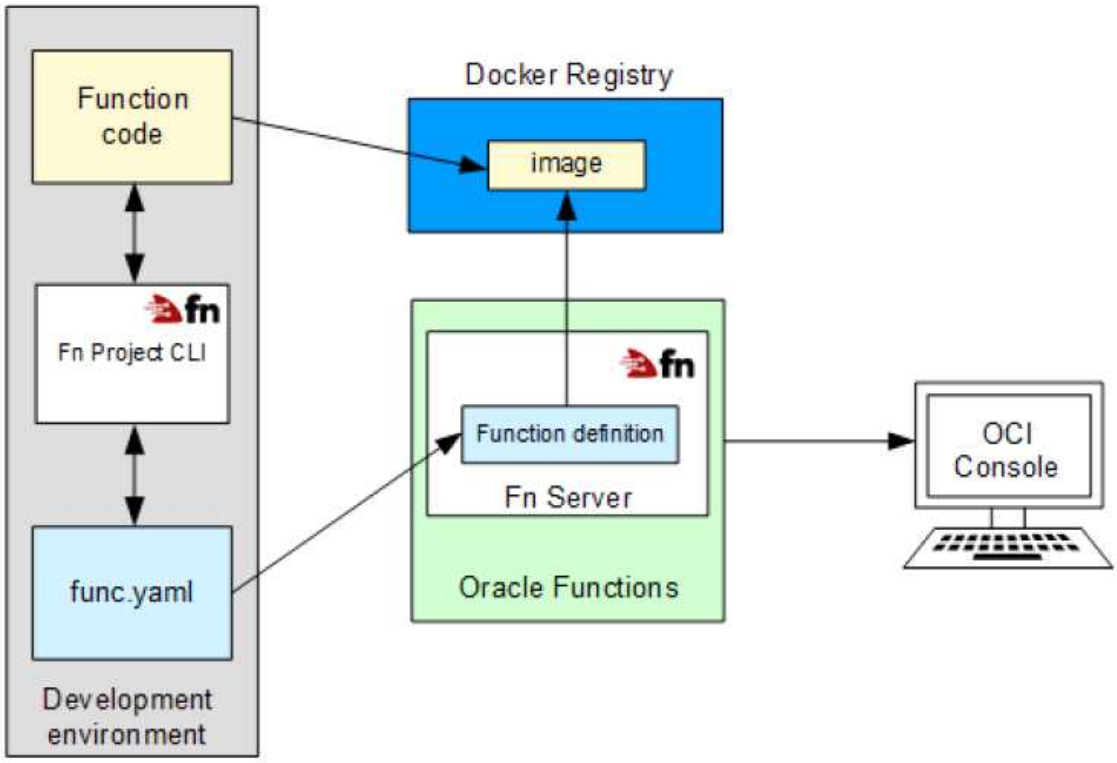

Oracle Cloud Infrastructure Container Engine for Kubernetes is a fully-managed, scalable, and highly available service that you can use to deploy your containerized applications to the cloud. Use Container Engine for Kubernetes (sometimes abbreviated to just OKE) when your development team wants to reliably build, deploy, and manage cloud-native applications. You specify the compute resources that your applications require, and Container Engine for Kubernetes provisions them on Oracle Cloud Infrastructure in an existing OCI tenancy.

Question: 5

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 it 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 themself 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.DeselO2.24 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.