

**100%** Money Back  
**Guarantee**

**Vendor:**Oracle

**Exam Code:**1Z0-997-20

**Exam Name:**Oracle Cloud Infrastructure 2020  
Architect Professional

**Version:**Demo

## QUESTION 1

An Oracle Cloud Infrastructure (OCI) Public Load Balancer's SSL certificate is expiring soon. You noticed the Load Balancer is configured with SSL Termination only. When the certificate expires, data traffic can be interrupted and security compromised.

What steps do you need to take to prevent this situation?

- A. Add the new SSL certificate to the Load Balancer, update backend servers to work with a new certificate and edit listeners so they can use the new certificate bundle.
- B. Add the new SSL certificate to the Load Balancer, update listeners and backend sets so they can use the new certificate bundle.
- C. Add the new SSL certificate to the Load Balancer and implement end to end SSL so it can encrypt the traffic from clients all the way to the backend servers.
- D. Add the new SSL certificate to the Load Balancer and update backend servers to use the new certificate bundle.
- E. Add the new SSL certificate to the Load Balancer and update listeners to use the new certificate bundle.

Correct Answer: C

Explanation: <https://docs.cloud.oracle.com/en-us/iaas/Content/Balance/Tasks/managingcertificates.htm>

---

## QUESTION 2

You are working as a solutions architect for an online retail store in Frankfurt which uses multiple compute instance VMs spread among three availability domains in the eu-frankfurt-1 region.

You noticed the website is having very high traffic, so you enabled autoscaling to solve one of your application but, you observed that one of the availability domains is not receiving any traffic.

What could be wrong in this situation?

- A. Autoscaling only works with single availability domains.
- B. You have to manually add all three availability domains to your load balancer configuration.
- C. Autoscaling can be enabled for multiple availability domains only in uk-london region.
- D. Autoscaling is using an Instance Pool configured to create instances in two availability Domains.
- E. You forgot to attach a load balancer to your instance pool configuration.

Correct Answer: D

Autoscaling lets you automatically adjust the number of Compute instances in an instance pool based on performance metrics such as CPU utilization. This helps you provide consistent performance for your end users during periods of high demand, and helps you reduce your costs during periods of low demand. You can associate a load balancer with an instance pool. If you do this, when you add an instance to the instance pool, the instance is automatically added to the load balancer's backend set. After the instance reaches a healthy state (the instance is listening on the configured port number), incoming traffic is automatically routed to the new instance. Instance pools let you provision and create

multiple Compute instances based off the same configuration, within the same region. By default, the instances in a pool are distributed across all fault Domains in a best-effort manner based on capacity. If capacity isn't available in one fault domain, the instances are placed in other fault domains to allow the instance pool to launch successfully. In a high availability scenario, you can require that the instances in a pool are evenly distributed across each of the fault domains that you specify. When sufficient capacity isn't available in one of the fault domains, the instance pool will not launch or scale successfully, and a work request for the instance pool will return an "out of capacity" error. To fix the capacity error, either wait for capacity to become available, or use the UpdateInstancePool operation to update the placement configuration (the availability domain and fault domain) for the instance pool. during create the instance pool you can select the location where you want to place the instances" In the Availability Domain list, select the availability domain to launch the instances in. If you want the instances in the pool to be placed evenly in one or more fault domains, select the Distribute instances evenly across selected fault domains check box. Then, select the fault domains to place the instances in.

---

### QUESTION 3

A cost conscious fashions design company which sells bags, clothes, and other luxury items has recently decided to move all of their on-premises infrastructure Oracle Cloud Infrastructure (OCI), One of their on-premises application is running on an NGINX server and the Oracle Database is running in a 2 node Oracle Real Application Clusters (RAC) configuration.

Based on cost considerations, what is an effective mechanism to migrate the customer application to OCI and set up regular automated backups?

- A. Launch a compute Instance and run a NGINX server to host the application. Deploy a 2 node VM DB Systems with oracle RAC enabled import the on premises database to OCI VM DB Systems using oracle Data Pump and then enable automatic backups.
- B. Launch a compute Instance and run an NGINX server to host the application. Deploy Exadata Quarter Rack, enable automatic backups and import the database using Oracle Data Pump.
- C. Launch a compute Instance for both the NGINX application server and the database server. Attach block volumes on the database server compute instance and enable backup policy to backup the block volumes.
- D. Launch a Compute instance and run a NGINX Server to host the application. Deploy a 2 node VM DB Systems with Oracle RAC enabled Import the on premises database to OCI VM DB Systems using data pump and then enable automatic backup- Also, enable Oracle Data Guard on the database server

Correct Answer: A

Based on cost considerations will exclude the Exadata. and there's no need for Data Guard Cost Estimator  
<https://www.oracle.com/cloud/cost-estimator.html>

Configuration Options		Pay As You Go	Monthly Flex	
▼	Database Cloud Service - OCI	\$17,190	\$11,460	🗑️
>	Database - OCI	\$17,190	\$11,460	🗑️
▼	Oracle Database Exadata Cloud Service	\$120,000	\$80,000	🗑️
>	Exadata	\$120,000	\$80,000	🗑️

#### QUESTION 4

An eCommerce company is running on Oracle Cloud Infrastructure (OCI) and many compute instances remain unused for the most part of the year except during Black Friday and Christmas. You suggest them to use OCI's

Autoscaling feature and present them a slide to showcase the features of Autoscaling.

Which option below is inaccurate in your presentation to the customer?

- A. A cooldown period between Autoscaling events lets the system stabilize at the updated level.
- B. When an instance pool scales in, instances are terminated in this order: the number of instances is balanced across Availability Domains, and then balanced across Fault Domains. Finally, within a Fault Domain, the newest instance is terminated first.
- C. Autoscaling relies on performance metrics such as CPU utilization that are collected by OCI Monitoring service to trigger an Autoscaling event.
- D. Autoscaling requires an instance pool as a pre-requisite so that it can automatically adjust the number of compute instances in an instance pool.

Correct Answer: B

#### 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 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.

Correct Answer: D

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.

---

### QUESTION 6

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

---

### QUESTION 7

You are working with a social media company as a solution architect. The media company wants to collect and analyze large amounts of data being generated from their websites and social media feeds to gain insights and continuously improve the user experience. In order to meet this requirement, you have developed a microservices application hosted on Oracle Container Engine for Kubernetes. The application will process the data and store the result to an Autonomous Data Warehouse (ADW) instance.

Which Oracle Cloud Infrastructure (OCI) service can you use to collect and process a large volume of unstructured data in real time?

- A. OCI Events
- B. OCI Streaming

C. OCI Resource Manager

D. OCI Notifications

Correct Answer: B

---

### QUESTION 8

After performing maintenance on an Oracle Linux compute instance the system is returned to a running state. You attempt to connect using SSH to do so. You decide to create an instance console connection to troubleshoot the issue.

Which three tasks would enable you to connect to the console connection and begin troubleshooting?

A. Use SSH to connect to the public IP address of the compute Instance and provide the console connection OCID as the username.

B. Edit the Linux boot menu to enable access to console.

C. Use SSH to connect to the service endpoint of the console connection service.

D. Reboot the compute instance using the Oracle Cloud Infrastructure (OCI) Management Console.

E. Upload an API signing key for console connection authentication.

F. Stop the compute Instance using the Oracle Cloud Infrastructure (OCI) Command Line interface (CLI).

Correct Answer: BCD

The Oracle Cloud Infrastructure Compute service provides console connections that enable you to remotely troubleshoot malfunctioning instances, such as:

An imported or customized image that does not complete a successful boot. A previously working instance that stops responding. the steps to connect to console and troubleshoot the OS Issue 1- Before you can connect to the serial console

you need to create the instance console connection.

Open the navigation menu. Under Core Infrastructure, go to Compute and click Instances.

Click the instance that you're interested in.

Under Resources, click Console Connections.

Click Create Console Connection.

Upload the public key (.pub) portion for the SSH key. You can browse to a public key file on your computer or paste your public key into the text box.

Click Create Console Connection.

When the console connection has been created and is available, the status changes to ACTIVE.

2- Connecting to the Serial Console

you can connect to the serial console by using a Secure Shell (SSH) connection to the service endpoint of the console connection service

Open the navigation menu. Under Core Infrastructure, go to Compute and click Instances.

Click the instance that you're interested in.

Under Resources, click Console Connections.

Click the Actions icon (three dots), and then click Copy Serial Console Connection for Linux/Mac.

Paste the connection string copied from the previous step to a terminal window on a Mac OS X or Linux system, and then press Enter to connect to the console. If you are not using the default SSH key or ssh-agent, you can modify the serial

console connection string to include the identity file flag, `-i`, to specify the SSH key to use. You must specify this for both the SSH

connection and the SSH ProxyCommand, as shown in the following line:

```
ssh -i // -o ProxyCommand=\\ssh -i // -W %h:%p -p 443...
```

Press Enter again to activate the console.

### 3- Troubleshooting Instances from Instance Console Connections To boot into maintenance mode

Reboot the instance from the Console.

When the reboot process starts, switch back to the terminal window, and you see Console messages start to appear in the window. As soon as you see the GRUB boot menu appear, use the up/down arrow

key to stop the automatic boot process, enabling you to use the boot menu. In the boot menu, highlight the top item in the menu, and type `e` to edit the boot entry. In edit mode, use the down arrow key to scroll down through the entries until

you reach the line that starts with either `linuxefi` for instances running Oracle Autonomous Linux 7.x or Oracle Linux 7.x, or `kernel` for instances running Oracle Linux 6.x.

At the end of that line, add the following:

```
init=/bin/bash
```

Reboot the instance from the terminal window by entering the keyboard shortcut `CTRL+X`.

---

## QUESTION 9

You have been asked to create a mobile application which will be used for submitting orders by users of a popular E-Commerce site. The application is built to work with Autonomous Transaction Processing - Serverless (ATP-S) database as the backend and HTML5 on Oracle Application Express as the front end. During the peak usage of the application you notice that the application response time is very slow. ATP-S database is deployed with 3 CPU cores and 1 TB of memory.

Which two options are expensive or impractical ways to improve the application response times?

- A. Identify the maximum memory capacity needed for peak times and scale the memory for the ATP-S database to that number. ATP-S will scale the memory down when not needed.
- B. Use the Machine Learning (ML) feature of the ATP-S database iteratively to tune the SQL queries used by the application.
- C. Scale up CPU core count and memory during peak times.
- D. Enable auto scaling for CPU cores on ATP-S database.
- E. Identify the maximum CPU capacity needed for peak times and scale the CPU core count for the ATP-S database to that number. ATP-S will scale the CPU core count down when not needed.

Correct Answer: CE

---

### QUESTION 10

You are helping a customer troubleshoot a problem. The customer has several Oracle Linux servers in a private subnet within a Virtual Cloud Network (VCN). The servers are configured to periodically communicate to the Internet to get security patches for applications installed on them.

The servers are unable to reach the Internet. An Internet Gateway has been deployed in the public subnet in the VCN and the appropriate routes are configured in the Route Table associated with the public subnet.

Based on cost considerations, which option will fix this issue?

- A. Create a Public Load Balancer in front of the servers and add the servers to the Backend Set of the Public Load Balancer.
- B. Create another Internet Gateway and configure it as route target for the private subnet.
- C. Implement a NAT instance in the public subnet of the VCN and configure the NAT instance as the route target for the private subnet.
- D. Create a NAT gateway in the VCN and configure the NAT gateway as the route target for the private subnet.

Correct Answer: D

---

### QUESTION 11

Which of the following is NOT a good use case for the volume backup feature of the Oracle Cloud Infrastructure Block Volume service?

- A. Support business continuity requirements of reducing the risk of outages or data mutation over time.
- B. Meet compliance and regulatory requirements for data to remain unchanged over time, so that it can be retrieved for audit purposes.
- C. Rapidly duplicate an environment in seconds to test configuration changes without impacting your production environment.
- D. Retain a copy of data in a volume, so that you can duplicate an environment later or preserve the data for future use.

Correct Answer: C

---

### QUESTION 12

You have been asked to implement a bespoke financial application in Oracle Cloud Infrastructure using virtual machine instances controlled by Autoscaling across multiple Availability Domains. The application stores transaction logs, intermediate transaction data, and audit data and needs to store this on a persistent, durable data store accessible from all of the application servers. The application requires the file system to be mounted in the /audit folder on the Linux file system. The system needs to tolerate the failure of two or more Fault Domains and still maintain data integrity. The solution should be as low maintenance as possible.

What storage architecture should you suggest?

- A. Use locally attached NVMe instances and configure RAID 0 replication between servers.
- B. Implement a single instance and install an NFS server, configure and create an NFS share, and mount this as /audit on the application instances.
- C. Store the data on Oracle Object Storage mounted at the /audit mount point on all the Linux instances using the default mount options.
- D. Use File Storage Service(FSS). Configure FSS to operate from all Availability Domains the application servers operate in and mount the file system in the /audit folder.

Correct Answer: D