



Welcome to download the Newest CertificationsTime 1Z0-1067 dumps
<https://certificationstime.com/updated/1z0-1067-exam-dumps-pdf/>

Exam Questions 1Z0-1067

Oracle Cloud Infrastructure 2019 Cloud Operations Associate

<https://certificationstime.com/>



<https://certificationstime.com/updated/1z0-1067-exam-dumps-pdf/>



Version: 6.2

Question: 1

You are asked to deploy a new application that has been designed to scale horizontally. The business stakeholders have asked that the application be deployed in us-phoenix-1.

Normal usage requires 2 OCPUs. You expect to have few spikes during the week, that will require up to 4 OCPUs, and a major usage uptick at the end of each month that will require 8 OCPUs.

What is the most cost-effective approach to implement a highly available and scalable solution?

- A. Create an instance pool with a VM.Standard2.2 shape instance configuration. Setup the autoscaling configuration to use 2 availability domains and have a minimum of 2 instances, to handle the weekly spikes, and a maximum of 4 instances.
- B. Create an instance with 1 OCPU shape. Use a CLI script to clone it when more resources are needed.
- C. Create an instance pool with a VM.Standard2.1 shape instance configuration. Setup the autoscaling configuration to use 2 availability domains and have a minimum of 2 instances and a maximum of 8 instances.
- D. Create an instance with 1 OCPU shape. Use the Resize Instance action to scale up to a larger shape when more resources are needed.

Answer: A

Explanation:

Instance pools let you provision and create multiple Compute instances based off the same instance configuration, within the same region. They also enable integration with other services, such as the Load Balancing service and IAM service, making it easier to manage groups of instances.

You create an instance pool using an existing instance configuration.

You can automatically adjust the number of instances in an instance pool based on performance metrics such as CPU utilization.

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.

<https://docs.cloud.oracle.com/en-us/iaas/Content/Compute/Tasks/creatinginstancepool.htm>

<https://blogs.oracle.com/cloud-infrastructure/autoscaling-a-load-balanced-web-application>

Question: 2

As the operations administrator for your company's Oracle Cloud Infrastructure (OCI), you have been entrusted the task of ensuring that data being accessed by the application is encrypted.

Your application portfolio includes both Virtual Machine (VM) and Bare Metal (BM) database systems.

Which method should you use to achieve encryption of data in-transit?



- A. Configure backup encryption for RMAN backup sets before transferring data
- B. Native Oracle Net Services encryption and integrity capabilities
- C. Key Store/Wallet service for on the fly encryption of data in transit
- D. Data is encrypted at rest using TDE and no additional encryption is needed

Answer: B

Explanation:

In Oracle Database Cloud Service databases, data security is provided for data in transit and data at rest. Security of data in transit is achieved through network encryption. Security of data at rest is achieved through encryption of data stored in database data files and backups.

To secure connections to your Oracle Database Cloud Service databases, you can use native Oracle Net Services encryption and integrity capabilities.

Encryption of network data provides data privacy so that unauthorized parties are not able to view data as it passes over the network. In addition, integrity algorithms protect against data modification and illegitimate replay.

Question: 3

You have deployed a three-tier web application inside an Oracle Cloud Infrastructure (OCI) VCN with a CIDR block of 10.0.0.0/28. You initially deploy three web servers (VM.Standard2.2), two application servers (VM.Standard2.4), and two servers (VM.Standard2.8) running Oracle database.

The web, application and database servers are deployed across two availability domains in the us-ashburn-1 region.

You also deployed a Public Load Balancer in front of the two web servers. The web traffic gradually increases in the first few days following the deployment, so you attempt to double the number of instances in each tier of the application to handle the new load. Unfortunately, some of these new instances fail to launch.

Your tenancy comes with the following set of predefined services limits for the availability domain and compartment where the application is deployed.

DESCRIPTION	LIMIT NAME	SERVICE LIMIT	USAGE	AVAILABLE
VM.Standard2.2	vm-standard2-2-count	30	0	30
VM.Standard2.4	vm-standard2-4-count	30	0	30
VM.Standard2.8	vm-standard2-8-count	10	0	10

What is a possible reason for this deployment to fail?

- A. You do not have enough private IP addresses left to launch all of the new compute instances.
- B. You do not have sufficient public IP addresses required by the web, application and database servers.
- C. You do not have sufficient quotas for number of VM.Standard2.2, VM.Standard2.4 and VM.Standard2.8 shapes in the Production compartment in the us-ashburn-1 region.
- D. You do not have sufficient quotas for number of VM.Standard2.2, VM.Standard2.4 and VM.Standard2.8 shapes in each availability domain in the us-ashburn-1 region.



Answer: A

Explanation:

Each subnet in a VCN consists of a contiguous range of IPv4 addresses that do not overlap with other subnets in the VCN.

In our case the VPN is use 10.0.0.0/28 and by default The first two IPv4 addresses and the last in the subnet's CIDR are reserved by the Networking service. so will have 13 free IPs should be used in this VCN We have already 3 web + 2 app +2 DB (7 IPs used for instances). in addition to 2 Private IPs for Load balance. that will give us 4 IPs only available in the Subnet that not allow us to double the VMs

Question: 4

A subscriber of on Oracle Cloud Infrastructure (OCI) Notifications service topic complained about not receiving messages from the service. Which of the following options can help you debug this issue?

- A. If OCI Notifications service does not receive an acknowledgement from a subscription endpoint, the service tries to redeliver messages for up to two hours. Configure an alarm on the NumberOfNotificationFailed metric through the OCI Monitoring service to help debug the issue.
- B. If OCI Notifications service does not receive an acknowledgement from a subscription endpoint, the service drops the message. Confirm that the subscriber is always online to receive messages to help debug the issue.
- C. If OCI Notifications service does not receive an acknowledgement from a subscription endpoint, the service tries to redeliver messages for up to one day. Make sure that the subscriber is online at least once a day to help debug the Issue.
- D. If OCI Notifications service does not receive an acknowledgement from a subscription endpoint, check the NumberOfNotificationFailed metric through the OCI Monitoring service for failed messages. Copy these messages to an OCI Object Storage bucket. Make sure the subscriber has the required credentials to access this bucket to help debug the Issue

Answer: A

Explanation:

The Oracle Cloud Infrastructure Notifications service broadcasts messages to distributed components through a publish-subscribe pattern, delivering secure, highly reliable, low latency and durable messages for applications hosted on Oracle Cloud Infrastructure and externally. Use Notifications to get notified when event rules are triggered or alarms are breached, or to directly publish a message.

If Notifications doesn't receive an acknowledgement from a subscription endpoint, the service tries to redeliver messages for up to two hours. This situation can occur when the endpoint is offline.

You can configure an alarm on the NumberOfNotificationFailed metric through the Monitoring service

Question: 5

The boot volume on your Oracle Linux instance has run out of space. Your application has crashed due to a lack of swap space, forcing you to increase the size of the boot volume.

Which step should NOT be included in the process used to solve the issue?



- A. Resize the boot volume by specifying a larger value than the boot volume's current size.
- B. Create a RAID 0 configuration to extend the boot volume file system onto another block volume.
- C. Attach the resized boot volume to a second instance as a data volume; Extend the partition and grow the file system on the resized boot volume.
- D. Reattach the boot volume and restart the instance.
- E. Stop the instance and detach the boot volume.

Answer: B

Explanation:

The Oracle Cloud Infrastructure Block Volume service lets you expand the size of block volumes and boot volumes. You have three options to increase the size of your volumes:

- Expand an existing volume in place with offline resizing. See [Resizing a Volume Using the Console](#) for the steps to do this.
- Restore from a volume backup to a larger volume. See [Restoring a Backup to a New Volume](#) and [Restoring a Boot Volume](#).
- Clone an existing volume to a new, larger volume. See [Cloning a Volume](#) and [Cloning a Boot Volume](#).

For the Full Access Visit:

<https://certificationstime.com/updated/1z0-1067-exam-dumps-pdf/>