DUMPSBOSS.

Oracle Cloud Infrastructure 2019 Cloud Operations Associate

Oracle 1z0-1067

Version Demo

Total Demo Questions: 10

Total Premium Questions: 73

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

QUESTION NO: 1

Which two statements accurately describe Ansible Modules for Oracle Cloud Infrastructure (OCI)?

A. OCI Ansible Modules represent discrete provisioning tasks or operations that you can not invoke individually from the command line, or else run individually or In sequence from a playbook.

B. OCI Ansible Modules are units of organization that allows you to abstract configuration, orchestration, and provisioning tasks into roles that you can save and share among playbooks and other users.

C. OCI Ansible Modules represent discrete provisioning tasks or operations that you can invoke individually from the command line, or else run Individually or in sequence from a playbook.

D. OCI Ansible Modules enable orchestrating, provisioning, and configuration management tasks on Oracle Cloud Infrastructure.

E. OCI Ansible Modules is not able to provide you state control of resources.

ANSWER: A D

Explanation:

Oracle supports the use of Ansible for cloud infrastructure provisioning, orchestration, and configuration management. Ansible allows you to automate configuring and provisioning your cloud infrastructure, deploying and updating software assets, and orchestrating your complex operational processes.

What enables orchestrating, provisioning, and configuration management tasks are the Ansible modules for Oracle Cloud Infrastructure. Ansible provides a library of these Ansible modules "out of the box" for managing common tasks, and libraries of custom modules from cloud providers like AWS and Azure. Oracle also provides a library of Ansible cloud modules that support provisioning and managing Oracle Cloud Infrastructure service

Ansible Modules represent discrete provisioning tasks or operations that you can invoke individually from the command line, or else run individually or in sequence from a playbook

Ansible roles are units of organization that allows you to abstract configuration, orchestration, and provisioning tasks into roles that you can save and share among playbooks and other users, and that are useful for organizing functionality in playbooks

https://docs.cloud.oracle.com/en-us/iaas/Content/API/SDKDocs/ansible.htm

QUESTION NO: 2

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-phoenlx-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 NO: 3

You have created the following JSON file to specify a lifecycle policy for one of your object storage buckets:

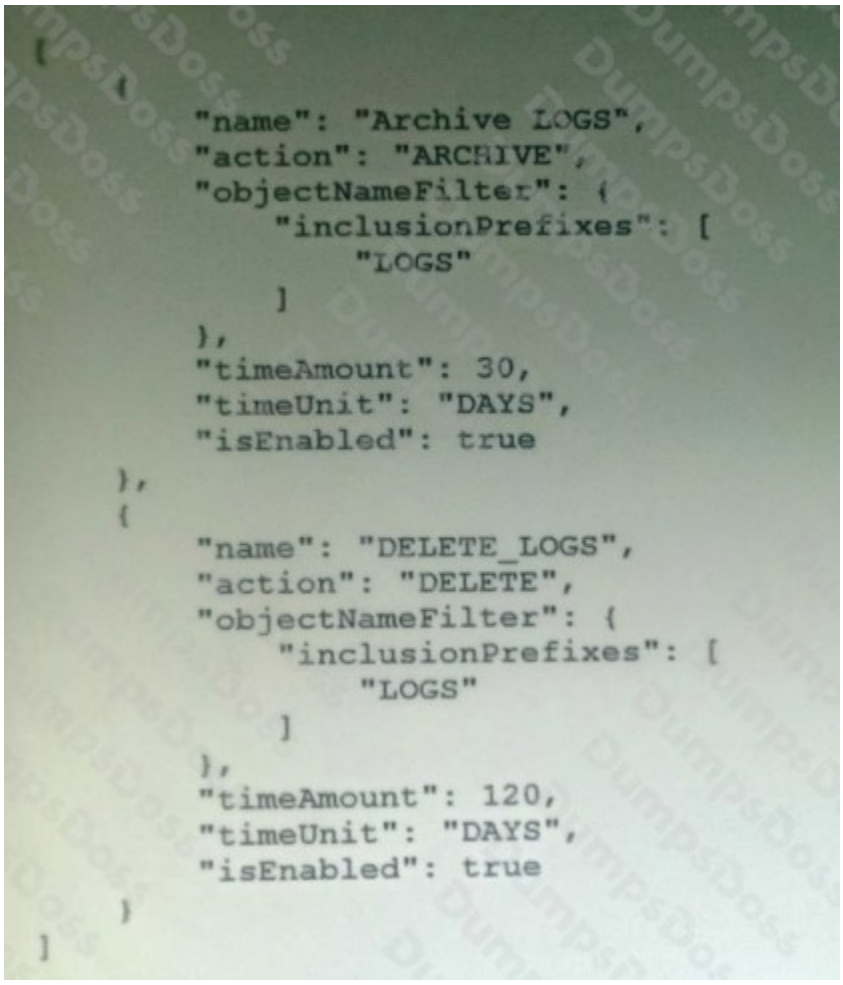

```
[
    {
        "name": "Archive LOGS",
        "action": "ARCHIVE",
        "objectNameFilter": {
            "inclusionPrefixes": [
                "LOGS"
            ]
        },
        "timeAmount": 30,
        "timeUnit": "DAYS",
        "isEnabled": true
    },
    {
        "name": "DELETE_LOGS",
        "action": "DELETE",
        "objectNameFilter": {
            "inclusionPrefixes": [
                "LOGS"
            ]
        },
        "timeAmount": 120,
        "timeUnit": "DAYS",
        "isEnabled": true
    }
]
```

How will this policy affect the objects that are stored in the bucket?

A. Objects containing the name prefix LOGS will be automatically migrated from standard Storage to Archive storage 30 days after the creation date. The objects will be deleted 120 days after creation.

B. Objects containing the name prefix LOGS will automatically be migrated from standard Storage to Archive storage 30 days after the creation date. The objects will be migrated back to standard Storage 120 days after creation.

C. The objects with prefix "LOGS" will be deleted 30 days after creation date.

D. Objects with the prefix "LOGS" will be retained for 120 days and then deleted permanently.

ANSWER: A

Explanation:

Using Object Lifecycle Management

Object Lifecycle Management lets you automatically manage the archiving and deletion of objects. By using Object Lifecycle Management to manage your Object Storage and Archive Storage data, you can reduce your storage costs and the amount of time you spend managing data.

Object Lifecycle Management works by defining rules that instruct Object Storage to archive or delete objects on your behalf within a given bucket. A bucket's lifecycle rules are collectively known as an object lifecycle policy.

This lifecycle policy archives objects after 30 days and deletes them after 120 days. for objects containing the name prefix LOGS

https://docs.cloud.oracle.com/en-us/iaas/Content/Identity/Reference/objectstoragepolicyreference.htm

QUESTION NO: 4

Your company will undergo a security audit in one week. Your manager has asked you to download and review recent logs from an Object Storage bucket. The current log archive file is approximately 19 GB In size.

Which command would you run to download the archive file as quickly as possible? A)

oci os object get -ns my-namespace -bn my-bucket --name my-large-object --multipart-download-threshold 2000 --part-size 120

B)

oci os object get -ns my-namespace -bn my-bucket --name my-large-object --multipart-download-threshold 2000 --part-size 128

C)

oci os object put -ns my-namespace -bn my-bucket --name my-large-object --multipart-download-threshold 20000 --part-size 128

D)

oci os object get -ns my-namespace -bn my-bucket --name my-large-object --multipart-download-threshold 20000 --part-size 128

A. Option A

B. Option B

C. Option C

D. Option D

ANSWER: B

Explanation:

Large files can be downloaded from Object Storage in multiple parts to speed up the download. You can configure the following options for the oci os object get command:

--multipart-download-threshold lets you specify the size, in MiB at which an object should be downloaded in multiple parts. This size must be at least 128 MiB.

--part-size, in MiB, to use for a download part. This gives you the flexibility to use more (smaller size) or fewer (larger size) parts as appropriate for your requirements. For example, compute power and network bandwidth. The default minimum part size is 120 MiB.

--parallel-download-count lets you specify how many parts are downloaded at the same time. A higher value may improve times but consume more system resources and network bandwidth. The default value is 10.

The following example shows the command to download any object with a size greater than 500 MiB. The object is downloaded in 128 MiB parts

oci os object get -ns my-namespace -bn my-bucket --name my-large-object --multipart-download-threshold 500 --part-size 128

--multipart-download-threshold [integer range]

Objects larger than this size (in MiB) will be downloaded in multiple parts. The minimum allowable threshold is 128 MiB.

https://docs.cloud.oracle.com/en-us/iaas/tools/oci-cli/2.9.1/oci_cli_docs/cmdref/os/object/get.html

QUESTION NO: 5

Which three statements ate true about Object Storage data security and encryption In Oracle Cloud Infrastructure (OCI)?

A. OCI Key Management is used by default to provide data security.

B. Client-side encryption is managed by the customer.

C. A VPN connection to OCI is required to ensure secure data transfer to an object storage bucket.

D. All traffic to and from Object Storage service is encrypted using TLS.

E. Server side encryption uses per-object keys which are managed by Oracle.

ANSWER: B D E

Explanation:

All data in Object Storage is encrypted at rest by using AES-256. Encryption is on by default and cannot be turned off. Each object is encrypted with its encryption key, and the object encryption keys are encrypted with a master encryption key. In addition, customers can use client-side encryption to encrypt objects with their encryption keys before storing them in Object Storage buckets. An available option for customers is to use the Amazon S3 Compatibility API, along with client-side object encryption support available in AWS SDK for Java.

Data in transit between customer clients (for example, SDKs and CLIs) and Object Storage public endpoints is encrypted with TLS 1.2 by default. FastConnect public peering allows on-premises access to Object Storage to go over a private network, rather than the public internet.

Oracle Cloud Infrastructure Key Management is a managed service the enables you, the customer, to manage and control AES symmetric keys used to encrypt your data-at-rest. Keys are stored in a FIPS 140-2, Level

3- certified, Hardware Security Module (HSM) that is durable and highly available. The Key Management service is integrated with many Oracle Cloud Infrastructure services, including Block Volumes, File Storage, Oracle Container Engine for Kubernetes, and Object Storage.

Use the Key Management service if you need to store your Master Encryption Keys in an HSM to meet governance and regulatory compliance requirements or when you want more control over the cryptoperiod of the encryption keys used for your data.

When you store your data with Oracle Cloud Infrastructure Block Volumes, File Storage Service, and Object Storage and don't use Key Management, your data is protected using encryption keys that are securely stored and controlled by Oracle.

QUESTION NO: 6

You are working as a Cloud Operations Administrator for your company. They have different Oracle Cloud Infrastructure (OCI) tenancies for development and production workloads. Each tenancy has resources in two regions - uk-london-1 and eu-frankfurt-1. You are asked to manage all resources and to automate all the tasks using OCI Command Line Interface (CLI).

Which is the most efficient method to manage multiple environments using OCI CLI?

A. Create environment variables for the sets of credentials that align to each combination of tenancy, region, and environment.

B. Use OCI CLI profiles to create multiple set of credentials in your config file, and reference the appropriate profile at runtime.

C. Use different bash terminals for each environment.

D. Run OCI setup config to create new credentials for each environment every time you want to access the environment.

ANSWER: B

Explanation:

The Oracle Cloud Infrastructure CLI configuration file can contain several profiles. and you can create multiple profiles with different values, then you can specify which profile to load.

Example Configuration [DEFAULT]

user=ocid1.user.oc1.. fingerprint= key_file=~/.oci/oci_api_key.pem tenancy=ocid1.tenancy.oc1.. region=us-ashburn-1

[ADMIN_USER]

user=ocid1.user.oc1.. fingerprint= key_file=keys/admin_key.pem pass_phrase=

https://docs.cloud.oracle.com/en-us/iaas/Content/API/Concepts/sdkconfig.htm

The Oracle Cloud Infrastructure CLI supports the use of environment variables to specify default values for some options and allows you to set environment variables to provide certain information. but the CLI requires a configuration file, See CLI Environment Variables for more information.

QUESTION NO: 7

You have created a public subnet in a VCN, and your public subnet has a Route Table, a Security List, and an Internet Gateway. However, none of the compute instances can connect to the Internet.

Which two are possible reasons for the connectivity issue? (Choose two.)

A. The Route Table has no default route for routing traffic to the Internet Gateway

B. There is no stateful ingress rule in the Security List associated with the public subnet

C. There is no Dynamic Routing Gateway (DRG) associated with the VCN

D. There is no stateful egress rule in the Security List associated with the public subnet

ANSWER: A D

Explanation:

An internet gateway as an optional virtual router that connects the edge of the VCN with the internet. To use the gateway, the hosts on both ends of the connection must have public IP addresses for routing. Connections that originate in your VCN and are destined for a public IP address (either inside or outside the VCN) go through the internet gateway. Connections that originate outside the VCN and are destined for a public IP address inside the VCN go through the internet gateway.

Working with Internet Gateways

You create an internet gateway in the context of a specific VCN. In other words, the internet gateway is automatically attached to a VCN. However, you can disable and re-enable the internet gateway at any time. Compare this with a dynamic routing gateway (DRG), which you create as a standalone object that you then attach to a particular VCN. DRGs use a different model because they're intended to be modular building blocks for privately connecting VCNs to your on-premises network.

For traffic to flow between a subnet and an internet gateway, you must create a route rule accordingly in the subnet's route table (for example, destination CIDR = 0.0.0.0/0 and target = internet gateway). If the internet gateway is disabled, that means no traffic will flow to or from the internet even if there's a route rule that enables that traffic. For more information, see Route Tables. For the purposes of access control, you must specify the compartment where you want the internet gateway to reside. If you're not sure which compartment to use, put the internet gateway in the same compartment as the cloud network. For more information, see Access Control.

You may optionally assign a friendly name to the internet gateway. It doesn't have to be unique, and you can change it later. Oracle automatically assigns the internet gateway a unique identifier called an Oracle Cloud ID (OCID). For more information, see Resource Identifiers.

To delete an internet gateway, it does not have to be disabled, but there must not be a route table that lists it as a target.

AS per compute instances can connect to the Internet so you use egress no ingress

QUESTION NO: 8

Which five are the required parameters to launch an instance in Oracle Cloud Infrastructure? (Choose five.)

A. private IPaddress

B. Virtual Cloud Network

C. host name

D. instance shape

E. image operating system

F. subnet

G. Availability Domain

ANSWER: B D E F G

Explanation:

https://docs.cloud.oracle.com/en-us/iaas/Content/Compute/Concepts/computeoverview.htm

QUESTION NO: 9

An organization wants to extend their existing on-premises data centers to the Oracle Cloud Infrastructure (OC1) us-phoenix-1 region. In order to achieve It, they have created an IPSec VPN connection between their Customer-Premises Equipment(CPE) and Dynamic Routing Gateway(DRG) on

How can you make this connection highly available (HA)?

A. Add another Dynamic Routing gateway In a different Availability Domain and create another IPSec VPN connection.

B. Add another Customer-Premises Equipment (CPE) and create second IPSec VPN connection with the same Dynamic Routing Gateway (DRG).

C. Create a NAT Gateway and route all traffic through a NAT Gateway, which is highly available component.

D. Add another Dynamic Routing Gateway in a different Availability Domain, and create another IPSec VPN connection with another Customer Premises Equipment (CPE).

ANSWER: B

Explanation:

IPSec VPN Best Practices

Configure all tunnels for every IPSec connection: Oracle deploys multiple IPSec headends for all your connections to provide high availability for your mission-critical workloads. Configuring all the available tunnels is a key part of the "Design for Failure" philosophy. (Exception: Cisco ASA policy-based configuration, which uses a single tunnel.)

Have redundant CPEs in your on-premises locations: Each of your sites that connects with IPSec to Oracle Cloud Infrastructure should have redundant CPE devices. You add each CPE to the Oracle Cloud Infrastructure Console and create a separate IPSec connection between your dynamic routing gateway

(DRG) and each CPE. For each IPSec connection, Oracle provisions two tunnels on geographically redundant IPSec headends. Oracle may use any tunnel that is "up" to send traffic back to your on-premises network. For more information, see Routing for the Oracle IPSec VPN.

Consider backup aggregate routes: If you have multiple sites connected via IPSec VPNs to Oracle Cloud Infrastructure, and those sites are connected to your on-premises backbone routers, consider configuring your IPSec connection routes with both the local site aggregate route as well as a default route.

Note that the DRG routes learned from the IPSec connections are only used by traffic you route from your VCN to your DRG. The default route will only be used by traffic sent to your DRG whose destination IP address does not match the more specific routes of any of your tunnels.

The following figure shows the basic layout of the IPSec VPN connection.

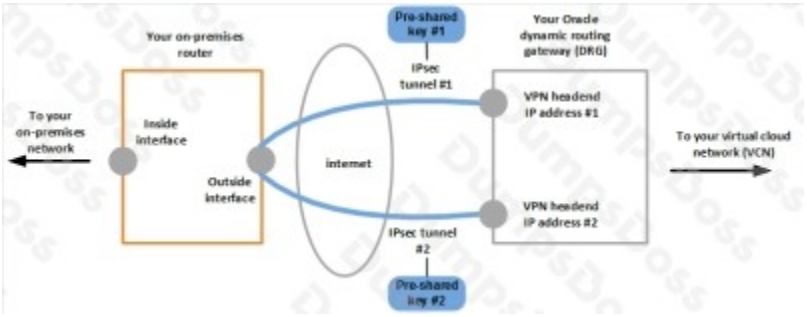

QUESTION NO: 10

Which two statements are true about Oracle Cloud Infrastructure Compute Service? (Choose two.)

A. You cannot launch a bare metal server in Oracle Cloud Infrastructure Compute Service

B. You can attach a block volume in an Availability Domain other than your compute instance

C. You can share custom images across tenancies and regions

D. You can launch a virtual or bare metal instance by using the same LaunchInstance API

ANSWER: C D

Explanation:

Regions and Availability Domains Volumes are only accessible to instances in the same availability domain . You cannot move a volume between availability domains or regions.