[UPDATED Oracle 1z0-997-22 Certification Exam Questions [Q63-Q79

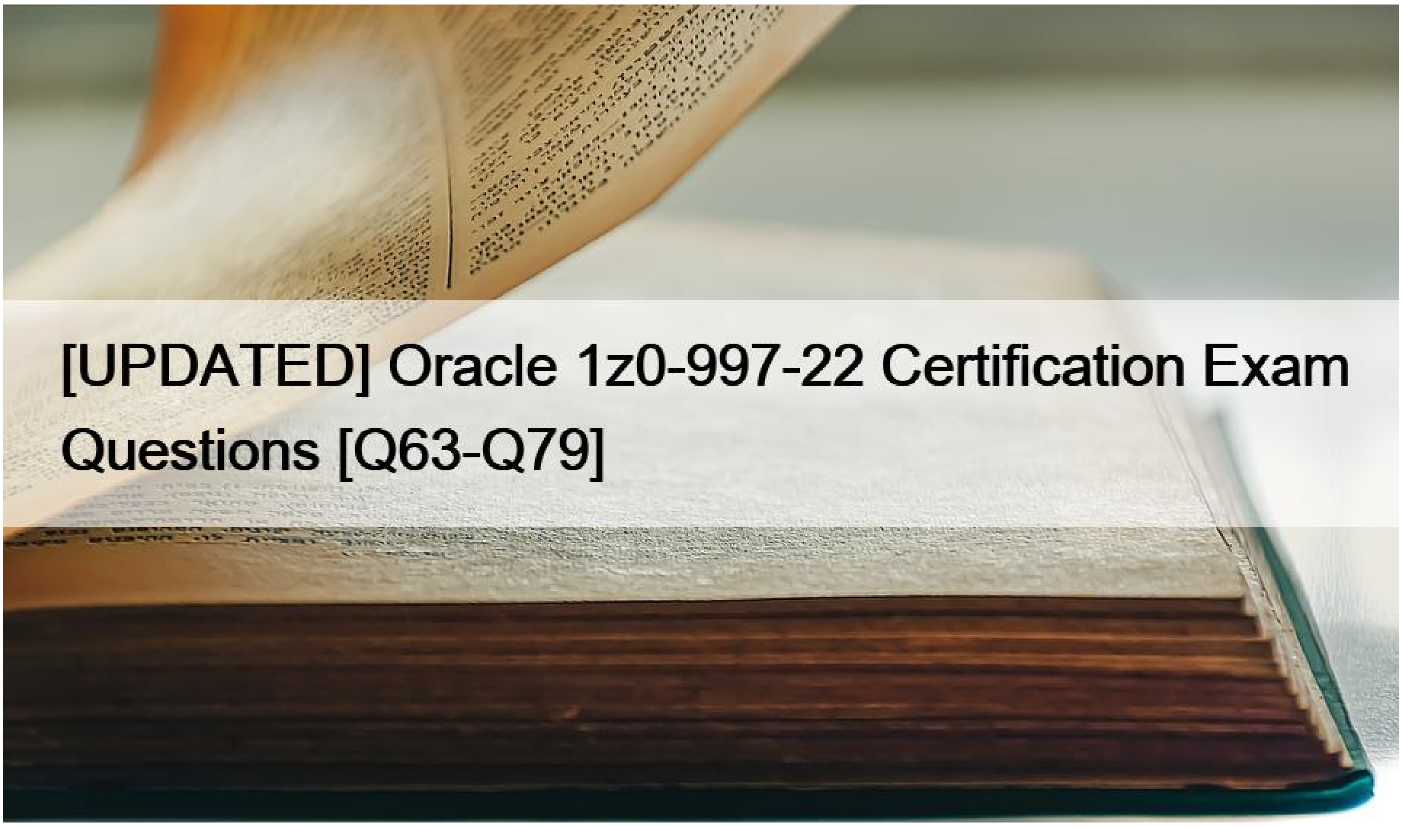

[UPDATED] Oracle 1z0-997-22 Certification Exam Questions
Quickly and Easily Pass Oracle Exam with 1z0-997-22 real Dumps

Oracle 1z0-997-22 exam is an excellent opportunity for IT professionals who are looking to enhance their knowledge and skills in the domain of Cloud Computing and Infrastructure. By passing this certification exam, you demonstrate your ability to design, implement, and maintain Oracle Cloud Infrastructure. Additionally, this certification enables you to enhance your professional value and demonstrate to potential employers that you possess the required knowledge and expertise required to perform well in the field of Cloud Computing and Infrastructure.

Oracle 1z0-997-22: Oracle Cloud Infrastructure 2022 Architect Professional exam is an excellent opportunity for professionals who want to validate their skills in designing and implementing cloud infrastructure solutions. 1z0-997-22 exam covers a broad range of topics and is designed to test the knowledge and abilities of the candidates in identifying and troubleshooting issues in the cloud infrastructure. Candidates can prepare for the exam by taking advantage of the various resources available online, including official Oracle Cloud Infrastructure documentation, training courses, and practice tests.

NO.63 You are creating an Oracle Cloud Infrastructure Dynamic Group. To determine the members of this group you are defining a set of matching rules.

Which of the following are the supported variables to define conditions in the matching rules? (Choose Two)
* instance.compartment.id -the OCID of the compartment where the instance resides.
* instance.tenancy.id -the OCID of the tenancy where the instance resides.
* tag.<tagnamespace>.<tagkey>.value -the tag namespace and tag key.
* iam.policy.id &#8211; the OCID of the IAM policy to apply to the group.

NO.64 You are responsible for migrating your on-premises legacy databases on 11.2.0.4 version to Autonomous Transaction Processing &#8211; Dedicated (ATP-D) in Oracle Cloud Infrastructure (OCI). As a solution architect, you need to plan your migration approach.

Which three options do you need to implement together to migrate your on-premises databases to OCI?
* Retain all legacy structures and unsupported features (e.g. legacy LOBs) in the on-premises databases for migration.
* Use Oracle Data Guard to keep on-premises database always active during migration.
* Launch Autonomous Transaction Processing &#8211; Dedicated database in OCI.
* Retain changes to Oracle shipped privileges, stored procedures or views in the on-premises databases.
* Convert on-premises databases to PDB, upgrade to 19c, and encrypt.
* Use Oracle GoldenGate replication to keep on-premises database online during migration.

NO.65 An online Stock trading application is deployed to multiple Availability Domains in the us phoenix-1 region. Considering the high volume of transactions that the trading application handles, the company has hired you to ensure that the data stored by the application available, and disaster resilient. In the event of failure, the Recovery lime Objective (UK)) must be less than 2 hours to meet regulator requirements.

Which Disaster Recovery strategy should be used to achieve the RTO requirement In the event of system failure?
* Configure hourly block volumes backups through the Storage Gateway service.
* Configure hourly block volumes backups using the Oracle Cloud Infrastructure (OCI) Command Line Interface (CLI)
* Store hourly block volumes backup to NVMe device under a compute instance and generate a custom Image every 5 minutes.
* Configure your application to use synchronous master slave data replication between Availability Domains.

You can use the CLI, REST APIs, or the SDKs to automate, script, and manage volume backups and their lifecycle.

Planning Your Backup

The primary use of backups is to support business continuity, disaster recovery, and long-term archiving requirements. When determining a backup schedule, your backup plan and goals should consider the following:

Frequency: How often you want to back up your data.

Recovery time: How long you can wait for a backup to be restored and accessible to the applications that use it. The time for a backup to complete varies on several factors, but it will generally take a few minutes or longer, depending on the size of the data being backed up and the amount of data that has changed since your last backup.

Number of stored backups: How many backups you need to keep available and the deletion schedule for those you no longer need. You can only create one backup at a time, so if a backup is underway, it will need to complete before you can create another one. For details about the number of backups you can store

NO.66 You are a solutions architect for a global health care company which has numerous data centers around the globe. Due to the ever growing data that your company is storing, you were Instructed to set up a durable, cost effective solution to archive you data from your existing on-premises tape based backup Infrastructure to Oracle Cloud Infrastructure (OCI).

What is the most-effective mechanism to Implement this requirement?

* Use the File Storage Service in OCI and copy the data from your existing tape based backup to the shared file system
* Setup an on premises OCI Storage Gateway which will back up your data to OCI Object Storage Archive tier.
* Setup an on premises OCI Storage Gateway which will back up your data to OCI object Storage Standard tier. Use Object Storage life cycle policy management to move any data older than 30 days from Standard to Archive tier.
* Setup an on-promises OCI Storage Gateway which will back up your data to OCI Object Storage Standard
* Setup fastConnect to connect your on premises network to your OCI VCN and use rsync tool to copy your data to OCI Object Storage Archive tier.

Oracle Cloud Infrastructure offers two distinct storage tiers for you to store your unstructured data. Use the Object Storage Standard tier for data to which you need fast, immediate, and frequent access. Use the Archive Storage service&#8217;s Archive tier for data that you access infrequently, but which must be preserved for long periods of time. Both storage tiers use the same manageable resources (for example, objects and buckets). The difference is that when you upload a file to Archive Storage, the object is immediately archived. Before you can access an archived object, you must first restore the object to the Standard tier.

you can use Storage Gateway to move files to Oracle Cloud Infrastructure Archive Storage as a cost effective backup solution. You can move individual files and compressed or uncompressed ZIP or TAR archives. Storing secondary copies of data is an ideal use case for Storage Gateway.

NO.67 You are working as a cloud consultant for a major media company. In the US and your client requested to consolidate all of their log streams, access logs, application logs, and security logs into a single system.

The client wants to analyze all of their logs In real-time based on heuristics and the result should be validated as well. This validation process requires going back to data samples extracted from the last 8 hours.

What approach should you take for this scenario?
* Create an auto scaling pool of syslog-enabled servers using compute instances which will store the logs In Object storage, then use map reduce jobs to extract logs from Object storage, and apply heuristics on the logs.
* Create a bare-metal instance big enough to host a syslog enabled server to process the logs and store logs on the locally attached NVMe SSDs for rapid retrieval of logs when needed.
* Set up an OCI Audit service and ingest all the API arils from Audit service pragmatically to a client side application to apply heuristics and save the result in an OCI Object storage.
* Stream all the logs and cloud events of Events service to Oracle Streaming Service. Build a client process that will apply heuristics on the logs and store them in an Object Storage.

The Oracle Cloud Infrastructure Streaming service provides a fully managed, scalable, and durable storage solution for ingesting continuous, high-volume streams of data that you can consume and process in real time. Streaming can be used for messaging, ingesting high-volume data such as application logs, operational telemetry, web click-stream data, or other use cases in which data is produced and processed continually and sequentially in a publish-subscribe messaging model.

Streaming Usage Scenarios

Here are some of the many possible uses for Streaming:

Metric and log ingestion: Use the Streaming service as an alternative for traditional file-scraping approaches to help make critical operational data more quickly available for indexing, analysis, and visualization.

Messaging: Use Streaming to decouple components of large systems. Streaming provides a pull/bufferbased communication model with sufficient capacity to flatten load spikes and the ability to feed multiple consumers with the same data independently.
Key-scoped ordering and guaranteed durability provide reliable primitives to implement various messaging patterns, while high throughput potential allows for such a system to scale well.

Web/Mobile activity data ingestion: Use Streaming for capturing activity from websites or mobile apps (such as page views, searches, or other actions users may take). This information can be used for realtime monitoring and analytics, as well as in data warehousing systems for offline processing and reporting.

Infrastructure and apps event processing: Use Streaming as a unified entry point for cloud components to report their life cycle events for audit, accounting, and related activities.

NO.68 You are tasked with backing up your data using Oracle Cloud Infrastructure Block Volume service.

When you are finalizing your block volume backup schedule, which of the following two are valid considerations for your backup plan? (Choose Two)
* Number of stored backups: How many backups you need to keep available and the deletion schedule for those you no longer need.
* Governance: Tagging of backups so you can capture backup related API calls through the Audit service.
* Frequency: How often you want to back up your data.
* Location: Determine the Object Store Bucket where the backups will be stored.
* Encryption: Whether to use your own key to encrypt your volume backups.

NO.69 You are designing the network infrastructure for an application consisting of a web server (server-1) and a Domain Name Server (server-2) running in two different subnets inside the same Virtual Cloud Network (VCN) in Oracle Cloud Infrastructure (OCI). You have a requirement where your end users will access server-1 from the internet and server-2 from your customer&#8217;s on-premises network. The on-premises network is connected to your VCN over a FastConnect virtual circuit.

How should you design your routing configuration to meet these requirements?
* Configure a single routing table with two set of rules: one that has route to internet via an Internet Gateway and another that propagates specific routes for the on-premises network via a Dynamic Routing Gateway. Don&#8217;t associate this routing table with any of the subnets in the VCN.
* Configure a single routing table with two set of rules: one that has route to internet via an Internet Gateway and another that propagate specific routes to the on-premises network via a Dynamic Routing Gateway. Associate the routing table with all the VCN subnets.
* Configure two routing tables: first one with a route to internet via an Internet gateway; associate this route table to the subnet containing server-1 .Configure the second route table to propagate specific routes to the on-premises network via a Dynamic Routing Gateway; associate this route table to subnet containing server-2.
* Configure two routing tables that have rules to route all traffic via a Dynamic Routing Gateway. Associate the two routing tables with all the VCN subnets.

NO.70 Which of the two options are true for an autonomous database in dedicated infrastructure deployment? (Choose two.)
* You can modify maintenance schedule of the AVM after provisioning, to match your organization maintenance schedules.
* The new resource model consists of autonomous exadata infrastructure, autonomous container database and autonomous database.
* Unlike autonomous database in shared infrastructure, you can customize the maintenance schedule of the autonomous databases in dedicated infrastructure in OCI public cloud.
* The new resource model consists of exadata infrastructure, autonomous Exadata VM cluster, autonomous container database.
* Network selection, License model and certificate management are resources configured at AVM level.

NO.71 You work for a bank as the lead Oracle Cloud Infrastructure architect. You designed a highly scalable solution for your company&#8217;s banking application. The architecture includes a load balancer, application servers with autoscaling configuration based on CPU utilization, and an Autonomous Database with Transaction Processing workload type running in a Virtual Cloud Network (VCN).

During the peak utilization period, the application users complain that the application runs slow.

What are two possible reasons for the application running slow at times? (Choose two.)
* The VCN does not have a Network Security Group configured to allow traffic from the load balancer to all the application servers in the backend set.
* Instance pool in autoscaling configuration for the application servers did not scale out due to compartment quota breach of the VM shapes used by the application servers.
* The load balancer is not configured correctly to send traffic to all the listeners of the application servers in the backend set.
* Instance pool in autoscaling configuration for the Autonomous Database did not scale out due to misconfigured scaling policy.
* Instance pool in autoscaling configuration for the application servers did not scale out due to service limit breach of the VM shapes used by the application servers.

NO.72 An Oracle Cloud Infrastructure (OCI) Public Load Balancer&#8217;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?
* 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.
* Add the new SSL certificate to the Load Balancer, update listeners and backend sets so they can use the new certificate bundle.
* 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.
* Add the new SSL certificate to the Load Balancer and update backend servers to use the new certificate bundle.
* Add the new SSL certificate to the Load Balancer and update listeners to use the new certificate bundle.
https://docs.cloud.oracle.com/en-us/iaas/Content/Balance/Tasks/managingcertificates.htm

NO.73 There are two compartments: Networks and Devlnstances

There are two groups: NetworkAdmins with a user named Nick, and Devs with a user named Dave The following IAM policies are being used:

*Allow group NetworkAdmins to manage virtual-network-family in compartment Networks

*Allow group NetworkAdmins to manage instance-family in compartment Networks

*Allow group Devs to use virtual-network-family in compartment Networks

*Allow group Devs to manage all-resources in compartment Devlnstances

Nick creates a VCN in Networks compartment. Dave creates a VCN in Devlnstances compartment.

Which of the following statements is INCORRECT?
* Dave launches instances in Devlnstances using the VCN in Networks compartment
* Nick cannot launch new instances in Devlnstances compartment
* Nick launches instances in Networks using VCN in Devlnstances compartment
* Dave cannot launch new instances in Networks compartment

NO.74 An organization has its IT infrastructure in a hybrid setup with an on-premises environment and an Oracle Cloud Infrastructure (OCI) Virtual Cloud Network (VCN) in the us-phonix-1 region. The on-premise applications communications with compute instances inside the VPN over a hardware VPN connection. They are looking to implement an Intrusion detected and Prevention (IDS/IPS) system for their OCI environment. This platform should have the ability to scale to thousands of compute of instances running inside the VCN.

How should they architect their solution on OCI to achieve this goal?
* Set up an OCI Private Load Balance! and configure IDS/IPS related health checks at TCP and/or HTTP level to inspect traffic
* Configure each host with an agent that collects all network traffic and sends that traffic to the IDS/IPS platform to inspection
* There Is no need to implement an IPS/IDS system as traffic coming over IPSec VPN tunnels Is already encrypt
* Configure autoscaling on a compute Instance pool and set vNIC to promiscuous mode to called traffic across the vcn and send it IDS/IPS platform for inspection.
in Transit routing through a private IP in the VCN you set up an instance in the VCN to act as a firewall or intrusion detection system to filter or inspect the traffic between the on-premises network and Oracle Services Network.

The Networking service lets you implement network security functions such as intrusion detection, application-level firewalls In fact, the IDS model can be host-based IDS (HIDS) or network-based IDS (NIDS). HIDS is installed at a host to periodically monitor specific system logs for patterns of intrusions. In contrast, an NIDS sniffs the traffic to analyze suspicious behaviors. A signature-based NIDS (SNIDS) examines the traffic for patterns of known intrusions. SNIDS can quickly and reliably diagnose the attacking techniques and security holes without generating an over-whelming number of false alarms because SNIDS relies on known signatures.

However, anomaly-based NIDS (ANIDS) detects unusual behaviors based on statistical methods. ANIDS could detect symptoms of attacks without specific knowledge of details. However, if the training data of the normal traffic are inadequate, ANIDS may generate a large number of false alarms.

NO.75 Your customer has gone through a recent departmental re structure. As part of this change, they are organizing their Oracle Cloud Infrastructure (OCI) compartment structure to align with the company&#8217;s new organizational structure.

They have made the following change:

Compartment x Is moved, and its parent compartment is now compartment c.

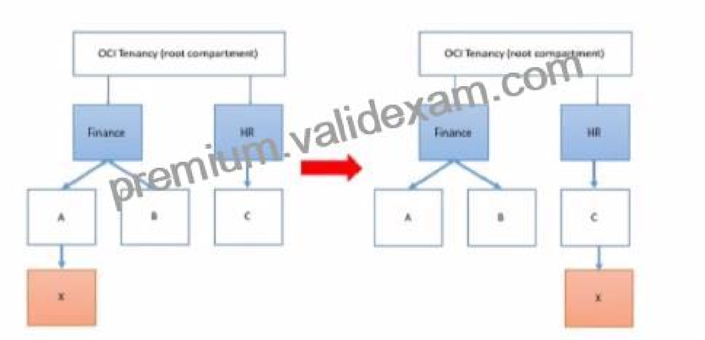
OCI Tenancy (root compartment)
Finance
HR
A
B
C
X
OCI Tenancy (root compartment)
Finance
HR
A
B
C
X

Policy defined in compartment A: Allow group networkadmins to manage subnets in compartment X Policy defined in root compartment: Allow group admins to read subnets in compartment Finance:A:X After you move the compartment, which two IAM policies would be required to ensure both groups retain the same permissions to compartment X that they had before? (Choose two.)
* Define a policy in the root compartment as follows: Allow group admins to manage subnets in compartment Finance:A:X
* Define a policy in compartment HR as follows: Allow group networkadmins to manage subnets in compartment C:X.
* Define a policy in the root compartment as follows: Allow group admins to read subnets in compartment HR:C:X
* Define a policy in compartment C as follows: Allow group networkadmins to read subnets in compartment X

NO.76 Your company will soon start moving critical systems Into Oracle Cloud Infrastructure (OCI) platform. These systems will reside in the us-phoenix-1and us-ashburn 1 regions. As part of the migration planning, you are reviewing the company&#8217;s existing security policies and written guidelines for the OCI platform usage within the company. you have to work with the company managed key.

Which two options ensure compliance with this policy?
* When you create a new compute instance through OCI console, you use the default options for &#8220;configure boot volume&#8221; to speed up the process to create this compute instance.
* When you create a new block volume through OCI console, select Encrypt using Key Management checkbox and use encryption keys generated and stored in OCI Key Management Service.
* When you create a new compute instance through OCI console, you use the default shape to speed up the process to create this compute instance.
* When you create a new OCI Object Storage bucket through OCI console, you need to choose &#8220;ENCRYPT USING CUSTOMER-MANAGED KEYS&#8221; option.
* You do not need to perform any additional actions because the OCI Block Volume service always encrypts all block volumes, boot volumes, and volume backups at rest by using the Advanced Encryption Standard (AES) algorithm with 256-bit encryption.
Block Volume Encryption

By default all volumes and their backups are encrypted using the Oracle-provided encryption keys. Each time a volume is cloned or restored from a backup the volume is assigned a new unique encryption key.

You have the option to encrypt all of your volumes and their backups using the keys that you own and manage using the Vault service.If you do not configure a volume to use the Vault service or you later unassign a key from the volume, the Block Volume service uses the Oracle-provided encryption key instead.

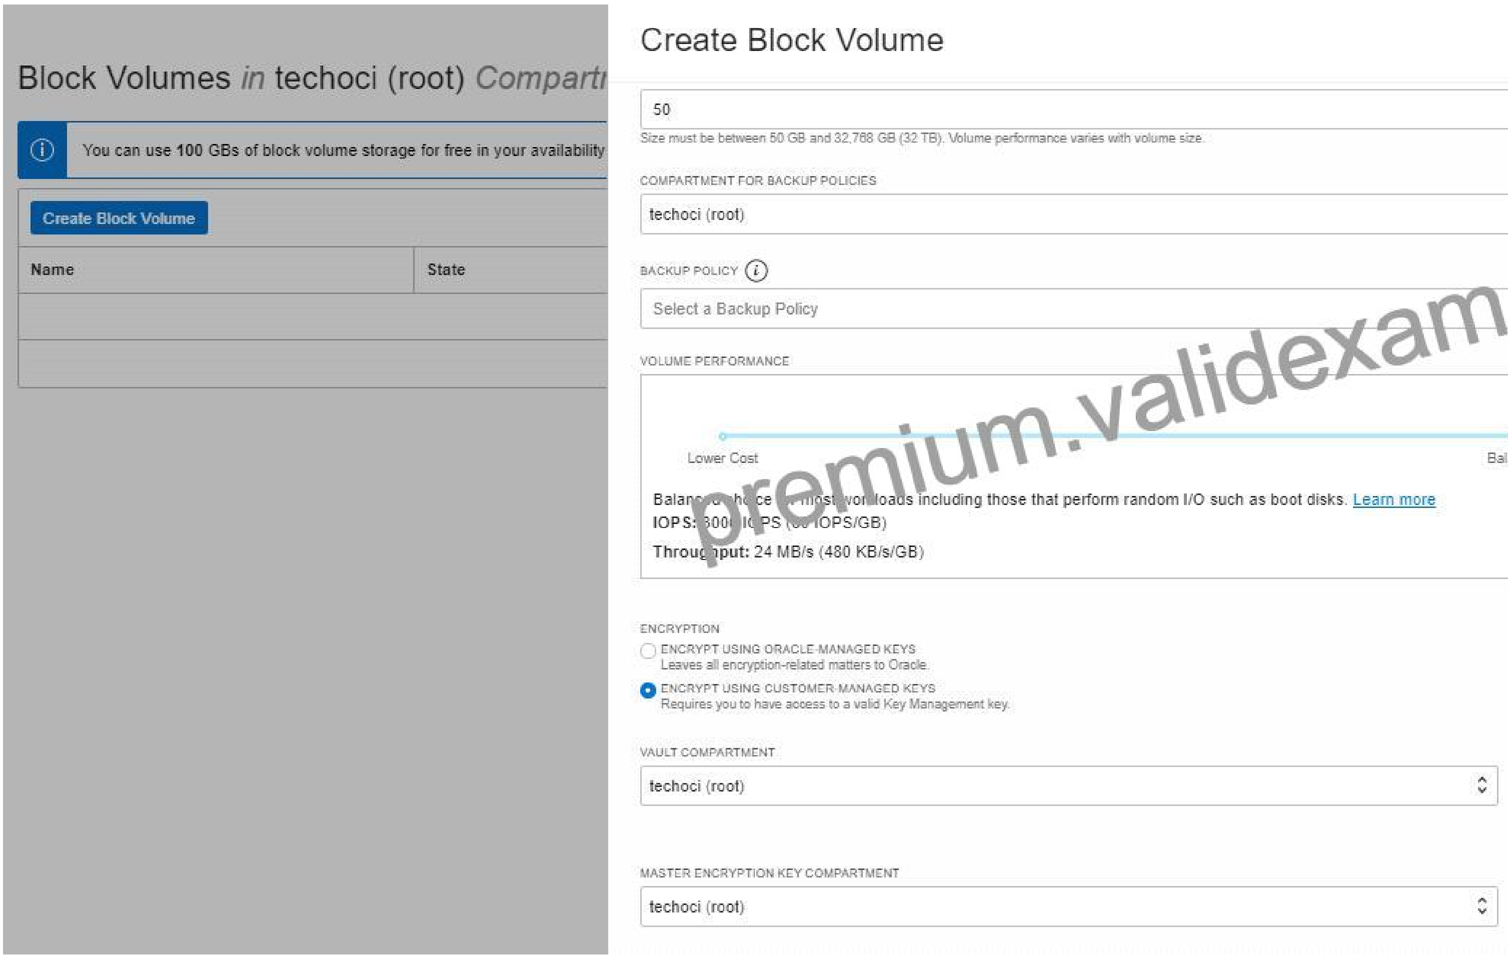

Block Volumes *in* techoci (root) *Compartı*

You can use **100** GBs of block volume storage for free in your availability

Name	State

Create Block Volume

50

Size must be between 50 GB and 32,768 GB (32 TB). Volume performance varies with volume size.

COMPARTMENT FOR BACKUP POLICIES: techoci (root)

BACKUP POLICY: Select a Backup Policy

VOLUME PERFORMANCE

Balanced choice for most workloads including those that perform random I/O such as boot disks. Learn more
IOPS: 3000 IOPS (60 IOPS/GB)
Throughput: 24 MB/s (480 KB/s/GB)

ENCRYPTION
○ ENCRYPT USING ORACLE-MANAGED KEYS
Leaves all encryption-related matters to Oracle.
● ENCRYPT USING CUSTOMER-MANAGED KEYS
Requires you to have access to a valid Key Management key.

VAULT COMPARTMENT: techoci (root)

MASTER ENCRYPTION KEY COMPARTMENT: techoci (root)

This applies to both encryption at-rest and in-transit encryption.

Object Storage Encryption

Object Storage employs 256-bit Advanced Encryption Standard (AES-256) to encrypt object data on the server. Each object is encrypted with its own data encryption key. Data encryption keys are always encrypted with a master encryption key that is assigned to the bucket. Encryption is enabled by default and cannot be turned off. By default, Oracle manages the master encryption key. However, you can optionally configure a bucket so that it&#8217;s assigned an Oracle Cloud Infrastructure Vault master encryption key that you control and rotate on your own schedule.

Encryption: Buckets are encrypted with keys managed by Oracle by default, but you can optionally encrypt the data in this bucket using your own Vault encryption key. To use Vault for your encryption needs, select Encrypt Using Customer-Managed Keys. Then, select the Vault Compartment and Vault that contain the master encryption key you want to use. Also select the Master Encryption Key Compartment and Master Encryption Key.

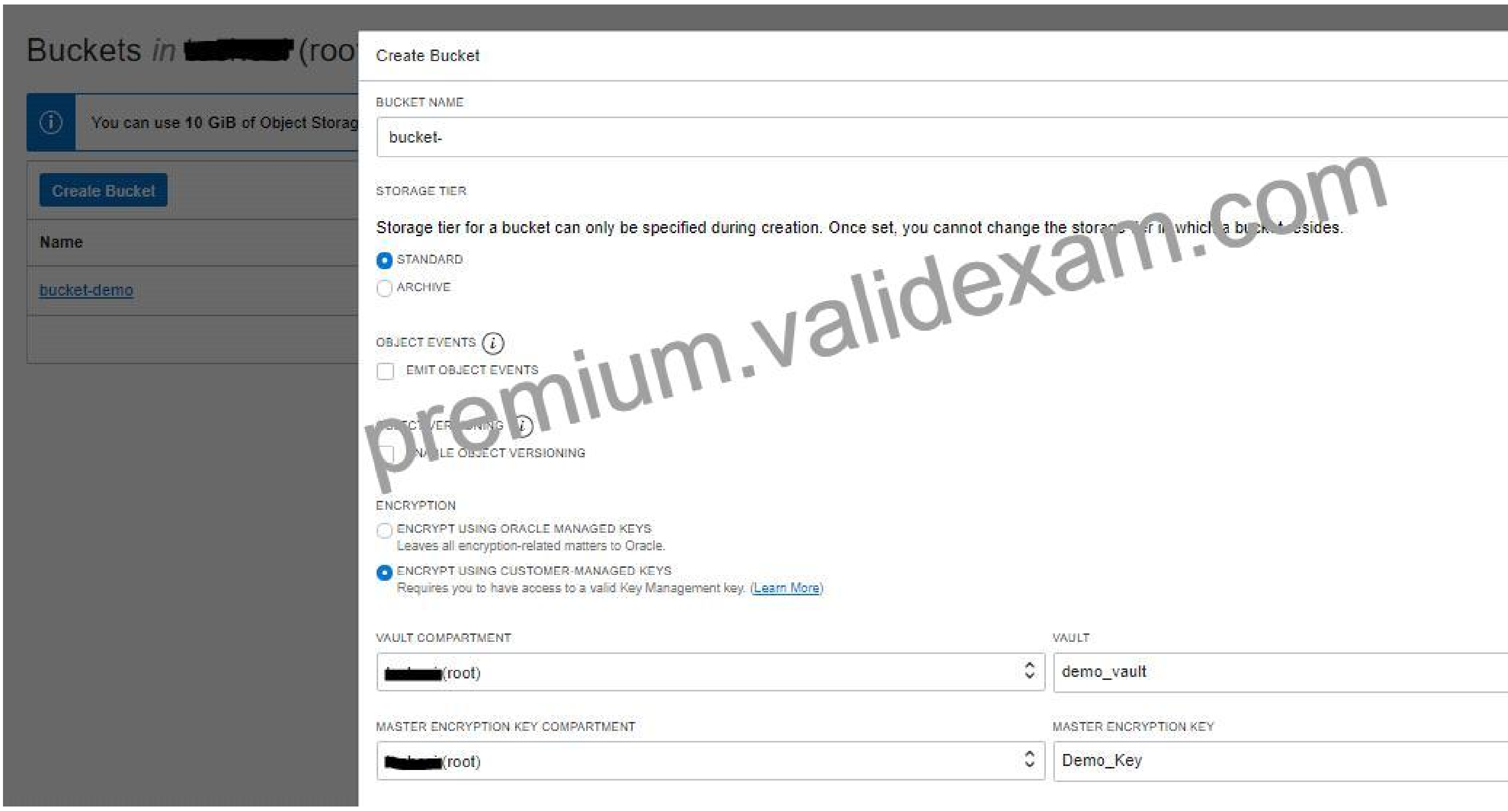

Buckets *in* [redacted] (roo

You can use 10 GiB of Object Storag

Create Bucket

BUCKET NAME

bucket-

STORAGE TIER

Storage tier for a bucket can only be specified during creation. Once set, you cannot change the storage tier in which a bucket resides.

● STANDARD

○ ARCHIVE

OBJECT EVENTS

□ EMIT OBJECT EVENTS

OBJECT VERSIONING

□ ENABLE OBJECT VERSIONING

ENCRYPTION

○ ENCRYPT USING ORACLE MANAGED KEYS
Leaves all encryption-related matters to Oracle.

● ENCRYPT USING CUSTOMER-MANAGED KEYS
Requires you to have access to a valid Key Management key. (Learn More)

VAULT COMPARTMENT: [redacted](root)

VAULT: demo_vault

MASTER ENCRYPTION KEY COMPARTMENT: [redacted](root)

MASTER ENCRYPTION KEY: Demo_Key

NO.77 After performing maintenance on an Oracle Linux compute instance the system is returned to a running state You attempt to connect using SSH t 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?
* Use SSH to connect to the public: IP address of the compute Instance and provide the console connection OCID as the username.
* edit the Linux boot menu to enable access to console.
* Use SSH to connect to the service endpoint of the console connection service
* Reboot the compute instance using the Oracle Cloud Infrastructure (OCI) Management Console

* Upload an API signing key for console connection authentication.
* Stop the compute Instance using the Oracle cloud Infrastructure (OCI) Command Line interface (CLI).

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&#8217;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&#8217;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 /<path>/<ssh_key> -o ProxyCommand=&#8217;ssh -i /<path>/<ssh_key> -W %h:%p -p 443&#8230;
```

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.

NO.78 You are working on the migration of the web application infrastructure of your company from on-premises to Oracle Cloud Infrastructure. You need to ensure that the DNS cache entries of external clients will not direct them to the on-premises infrastructure after switching to the new infrastructure.

Which of the following options will minimize this problem?
* Reduce the TTL of the DNS records after the switch.
* DNS changes propagate fast enough that it is not necessary to take any action.
* Increase the TTL of the DNS records before the switch.
* Increase the TTL of the DNS records after the switch.
* Reduce the TTL of the DNS records before the switch.

NO.79 You are working as a solution architect for a customer in Frankfurt, which uses multiple compute instance VMs spread among three Availability Domains in the Oracle Cloud Infrastructure (OCI) eu-frankfurt-1 region. The compute instances do not have public IP addresses and are running in private subnets inside a Virtual Cloud Network (VCN). You have set up OCI Autoscaling feature for the compute instances, but find out that instances cannot be auto scaled. You have enabled monitoring on the instances.

What could be wrong in this situation?
* You need to assign a reserved public IP address to the compute instances.
* You need to set up a Service Gateway to send metrics to the OCI Monitoring service.
* Autoscaling only works for instances with public IP addresses.
* Autoscaling only works with single availability domains.

Start your 1z0-997-22 Exam Questions Preparation: https://www.validexam.com/1z0-997-22-latest-dumps.html]