



1Z0-997-22^{Q&As}

Oracle Cloud Infrastructure 2022 Architect Professional

Pass Oracle 1Z0-997-22 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.passapply.com/1z0-997-22.html>

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Oracle
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





QUESTION 1

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?

- A. Use the File Storage Service in OCI and copy the data from your existing tape based backup to the shared file system
- B. Setup an on premises OCI Storage Gateway which will back up your data to OCI Object Storage Archive tier.
- C. 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.
- D. Setup an on-promises OCI Storage Gateway which will back up your data to OCI Object Storage Standard
- E. 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.

Correct Answer: B

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'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.

QUESTION 2

A startup company is looking for a solution for processing of data transmitted by the IOT devices fitted to transport vehicles that carry frozen foods. The data should be consumed and processed in real time. The processed data should be archived to OCI Object Storage bucket. and use Autonomous Data warehouse (ADW) to handle analytics.

Which architecture will help you meet this requirement?

- A. Use OCI Streaming Service to collect the incoming biometric data. Use an open source Hadoop cluster to analyze the data horn streaming service. Store the results to OCI Autonomous Data warehouse (ADW) to handle complex analytics
- B. Use OCI Streaming Service to collect the incoming biometric data. Use Oracle Functions to process the date and show the results on a real-time dashboard and store the results lo OCI Object Storage Store the data In OCI Autonomous Data warehouse (ADW) to handle analytics.
- C. Create an OCI Object Storage bucket to collect the incoming biometric data from the smart pet collar Fetch the data horn OC\ Object storage to OCI Autonomous Data Warehouse (ADW) every day and run analytics Jobs with it
- D. Launch an open source Hadoop cluster to collect the Incoming biometrics data Use an Open source Fluentd cluster



to analyze the- data me results to OCI Autonomous Transaction Processing (ADW)to handle complex analytics

Correct Answer: B

Real-time processing of high-volume streams of data

-OCI Streaming service provides a fully managed, scalable, durable storage option for continuous, highvolume streams of data that you can consume and process in real-time

-Use cases Log and Event data collection Web/Mobile activity data ingestion IoT Data streaming for processing and alerts Messaging: use streaming to decouple components of large systems

-Oracle managed service with REST APIs (Create, Put, Get, Delete)

-Integrated Monitoring

QUESTION 3

You work for a large bank where security and compliance are critical. As part of the security overview meeting, your company decided to minimize the installation of local tools on your laptop. You have been running Ansible and kubectl to spin up Oracle Container Engine for Kubernetes (OKE) clusters and deployed your application.

For authentication, you are using an Oracle Cloud Infrastructure (OCI) CLI config file that contains OCIDs, Fingerprint, and a locally stored PEM file. Your security team doesn't want you to store any local API key and certificate, or any other local tools.

Which two actions should you perform to spin up the OKE cluster and interact with it? (Choose two.)

A. Create a developer workstation on OCI. Install Ansible and kubectl on it. Use resource principal to authenticate against OCI API and create the OKE Cluster.

B. Develop your own code using OCI SDK to deploy the OKE cluster.

C. Work on OCI Cloud Shell to use built-in Ansible and kubectl to deploy the OKE cluster. Use OCI_CLI_AUTH=instance_obo_user environment variable to authenticate using built- in token.

D. Work on OCI Cloud Shell to use built-in Ansible and kubectl to deploy the OKE cluster. Bring in your own config file and certificate to authenticate against OCI API.

E. Create a developer workstation on OCI. Install Ansible and kubectl on it. Use instance principal to authenticate against OCI API and create the OKE Cluster.

Correct Answer: CE

https://docs.cloud.oracle.com/en-us/iaas/tools/oci-cli/2.12.4/oci_cli_docs/oci.html

QUESTION 4

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?

- A. Configure hourly block volumes backups through the Storage Gateway service.
- B. Configure hourly block volumes backups using the Oracle Cloud Infrastructure (OCI) Command Line Interface (CLI)
- C. Store hourly block volumes backup to NVMe device under a compute instance and generate a custom Image every 5 minutes.
- D. Configure your application to use synchronous master slave data replication between Availability Domains.

Correct Answer: B

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

QUESTION 5

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)

- A. instance.compartment.id -the OCID of the compartment where the instance resides.
- B. instance.tenancy.id -the OCID of the tenancy where the instance resides.
- C. tag...value -the tag namespace and tag key.
- D. iam.policy.id - the OCID of the IAM policy to apply to the group.

Correct Answer: AC

[1Z0-997-22 PDF Dumps](#)

[1Z0-997-22 VCE Dumps](#)

[1Z0-997-22 Braindumps](#)