

1Z0-997-21^{Q&As}

Oracle Cloud Infrastructure 2021 Architect Professional

Pass Oracle 1Z0-997-21 Exam with 100% Guarantee

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

<https://www.leads4pass.com/1z0-997-21.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 building a demo for a customer that showcases Oracle Cloud Infrastructure (OCI) Events service and Oracle Functions. You plan to create an event every time an image is uploaded to an OCI Object Storage bucket. You have also created a function that is listening to the event and processes the image for face recognition. Choose the two actions from below that are NOT required to run the demo successfully.

- A. You must specify an action type while creating an Event service and specify the function you want to trigger.
- B. Creating an event rule is not permitted for OCI Object storage.
- C. The function must be deployed only to Oracle Kubernetes Engine (OKE).
- D. You have to enable Object Storage buckets to emit events for state changes.
- E. You must deploy the function that does facial recognition for the demo to work.

Correct Answer: BC

QUESTION 2

You work for a large bank where your main application is a payment processing gateway API. You deployed the application on Oracle Container Engine for Kubernetes (OKE) and used API Gateway with several policies to control the access of the API endpoint. However, your customers are complaining about the unavailability of the API endpoint. Upon checking, you noticed that the Gateway URL is throwing Service Unavailable error. You need to check the backend latency and backend responses when this error started last night. What should you do to get this data? (Choose the best answer.)

- A. Check with the application owner and search the log file for the container to get the metrics from the log file.
- B. Go to Governance Menu and click on Audit to see the Audit log for the API Gateway. Filter it using Start and End date with a 503 response status.
- C. Go to Developer Services and click on API Gateway. Go to the detail page of the gateway and select Metrics. Change the Start and End time to filter the metrics.
- D. Go to Monitoring and click on Service Metrics. Choose the Metric Namespace as oci_apigateway. Change the Start and End time accordingly. Add a Dimension and select httpStatusCode: 503. Check the backend latency and backend responses metric.

Correct Answer: D

<https://medium.com/oracledevs/using-oci-monitoring-healthchecks-to-schedule-execution-of-serverlessfunctions-on-oracle-cloud-ef233f887a5>

QUESTION 3

An OCI Architect is working on a solution consisting of analysis of data from clinical trials of a pharmaceutical company. The data is being stored in OCI Autonomous Data Warehouse (ADW) having 8 CPU Cores and 70 TB of storage. The architect is planning to setup autoscaling to respond to dynamic changes in the workload. Which of the following needs to be considered while configuring auto scaling? Choose two

- A. Enabling auto scaling does not change the concurrency and parallelism settings
- B. Auto scaling also scales IO throughput linearly along with CPU
- C. The database memory SGA and PGA will not get affected by the changes in the number of CPUs during auto scaling
- D. The maximum CPU cores that will be automatically allocated for this database is 16 CPUs

Correct Answer: AB

Auto scaling is enabled by default when you create an Autonomous Database instance or you can use Scale Up/Down on the Oracle Cloud Infrastructure console to enable or disable auto scaling. With auto scaling enabled the database can use up to three times more CPU and IO resources than specified by the number of OCPUs currently shown in the Scale Up/Down dialog. When auto scaling is enabled, if your workload requires additional CPU and IO resources the database automatically uses the resources without any manual intervention required. Enabling auto scaling does not change the concurrency and parallelism settings for the predefined services IO throughput depends on the number of CPUs you provision and scales linearly with the number of CPUs.

QUESTION 4

Which of the following is NOT a good use case for the Oracle Cloud Infrastructure (OCI) Streaming service?

- A. Meeting compliance requirements for data to remain unchanged over a long time, so that it can be retrieved for audit purposes.
- B. Messaging with a pull-based communication model and the ability to feed multiple consumers with the same data independently.
- C. Ingesting metric and log data to help make critical operational data more quickly available for indexing, analysis, and visualization.
- D. Providing a unified entry point for cloud components to report their life cycle events for audit, accounting, and related activities.

Correct Answer: A

QUESTION 5

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 is available, and disaster resilient. In the event of failure, the Recovery Time Objective (RTO) 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

[1Z0-997-21 VCE Dumps](#)

[1Z0-997-21 Study Guide](#)

[1Z0-997-21 Exam Questions](#)