



1Z0-1072-21^{Q&As}

Oracle Cloud Infrastructure 2021 Architect Associate

Pass Oracle 1Z0-1072-21 Exam with 100% Guarantee

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

<https://www.pass4itsure.com/1z0-1072-21.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**

Which two options are true for Autonomous Transaction Processing (ATP) database? (Choose two.)

- A. You can add/remove Diskgroup in ATP
- B. You can scale storage up or down in ATP
- C. You can scale CPU up or down in ATP
- D. You can add more Pluggable Database for consolidating multiple databases in ATP
- E. You can add new ORACLE_HOME for bringing older versions of on-premises databases to ATP

Correct Answer: BC

You can scale up/down your Autonomous Database to scale both in terms of compute and storage only when needed, allows people to pay per use.

Oracle allows you to scale compute and storage independently, no need to do it together. these scaling activities fully online (no downtime required) in Details page Autonomous Database click Scale Up/Down.

Click on arrow to select a value for CPU Core Count or Storage (TB).

Or Select auto scaling to allow the system to automatically use up to three times more CPU and IO resources to meet workload demand, compared to the database operating with auto scaling disabled.

QUESTION 2

Which statement is true regarding Autonomous Transaction Processing (ATP)?

- A. A database name cannot be used concurrently for both an Autonomous Data Warehouse (ADW) and an ATP database
- B. After terminating a database, the database name is available for immediate reuse
- C. A maximum of 8 cores can be enabled for an ATP database
- D. A maximum of 2 TB of storage can be enabled for an ATP database

Correct Answer: A

The database name must be unique among all Autonomous Data Warehouses and Autonomous Databases in your tenancy in the same region.

 Provisioning failed because a database named ██████████ already exists in compartment adb_compartment. The name must be unique among all Autonomous Data Warehouses and Autonomous Databases in your tenancy in the same region. Specify a different database name and try again.



Terminating an Autonomous Transaction Processing database permanently deletes the instance and removes all automatic backups. You cannot recover a terminated database. The maximum number of CPUs and maximum storage capacity that can be provisioned in Oracle Autonomous Database in the current release up to 128 CPUs and 128TB can be provisioned from the cloud console. Customers requiring more resources need to call their Oracle account team

QUESTION 3

Which statement is true about Data Guard implementation in Oracle Cloud Infrastructure (OCI) bare metal and virtual machine database systems?

- A. Primary and standby databases must be in the same OCI region.
- B. Both database systems must be in the same compartment.
- C. Database systems need not be the same shape type (e.g, primary database can be a virtual machine, and standby database a bare metal shape, and vice versa).
- D. Primary and standby database versions and editions need not be identical.

Correct Answer: B

Reference: <https://docs.cloud.oracle.com/en-us/iaas/Content/Database/Tasks/exausingdataguard.htm>

QUESTION 4

Which of the following statements is true regarding Oracle Cloud Infrastructure Object Storage Pre-Authenticated Requests?

- A. It is not possible to create pre-authenticated requests for "archive" storage tier
- B. Changing the bucket visibility does not change existing pre-authenticated requests
- C. It is not possible to create pre-authenticated requests for the buckets, but only for the objects
- D. Pre-authenticated requests don't have an expiration

Correct Answer: B

Pre-authenticated requests provide a way to let users access a bucket or an object without having their own credentials, as long as the request creator has permissions to access those objects. For example, you can create a request that lets an operations support user upload backups to a bucket without owning API keys. Or, you can create a request that lets a business partner update shared data in a bucket without owning API keys. When you create a pre-authenticated request, a unique URL is generated. Anyone you provide this URL to can access the Object Storage resources identified in the pre-authenticated request, using standard HTTP tools like curl and wget. Understand the following scope and constraints regarding pre-authenticated requests: Users can't list bucket contents. You can create an unlimited number of pre-authenticated requests. There is no time limit to the expiration date that you can set. You can't edit a pre-authenticated request. If you want to change user access options in response to changing requirements, you must create a new pre-authenticated request. The target and actions for a pre-authenticated request are based on the creator's permissions. The request is not, however, bound to the creator's account login credentials. If the creator's login credentials change, a pre-authenticated request is not affected. You cannot delete a bucket that has a pre-authenticated request associated with that bucket or with an object in that bucket. Understand the following scope and constraints regarding public access: Changing the type of access is bi-directional. You can change a bucket's access from public to private or from private to public. Changing the type of access doesn't affect existing pre-authenticated



requests. Existing pre- authenticated requests still work.

QUESTION 5

You currently manage an e-commerce application that utilizes 25 identical compute resources to handle customer traffic. The stakeholders have asked you to create another 25 identical compute resources in order to deploy and test a new version of the software? What is the most efficient process to create 25

additional compute resources that are identical to the first 25?

- A. Create a custom image from 1 of the 25 servers. Use this custom image to provision 25 more servers
- B. Create a manual backup of each boot volume belonging to the 25 servers. Restore each backup to create 25 new boot volumes, from which you will provision 25 more servers
- C. Provision a new server and configure it to be identical to the first 25. Create a custom image from the new server, then use the custom image to provision 24 more servers
- D. Clone the boot volume of 1 of the 25 servers. Use the boot volume clone to provision 25 more servers

Correct Answer: A

[1Z0-1072-21 PDF Dumps](#)

[1Z0-1072-21 Practice Test](#)

[1Z0-1072-21 Exam Questions](#)