

100% Money Back
Guarantee

Vendor:Oracle

Exam Code:1Z0-931-22

Exam Name:Oracle Autonomous Database Cloud
2022 Professional

Version:Demo

QUESTION 1

In the Autonomous Database (ADB) Dedicated service, what does the fleet administrator use to control OCPU utilization?

- A. Oracle Machine Learning notebook
- B. Compartment quotas
- C. Resource Manager settings
- D. SQL Developer Web Console

Correct Answer: B

<https://docs.oracle.com/en-us/iaas/Content/Database/Concepts/adboverview.htm>

Compartment Quotas for Autonomous Databases

You can use compartment quotas to control how Autonomous Database OCPU and storage resources are allocated to Oracle Cloud Infrastructure compartments. You can use compartment quota policy statements to control OCPU and storage resources by both [workload type](#) and [Exadata infrastructure type](#). For example, you can allocate 10 Autonomous Transaction Processing OCPUs on shared Exadata infrastructure to a specific compartment. This would not affect the number of OCPUs available to Autonomous Data Warehouse databases, or databases using dedicated Exadata infrastructure. For more information on using compartment quotas, see [Compartment Quotas](#) and [Database Quotas](#).

QUESTION 2

Clone Sources

You can use a running database to create a clone. For databases running on shared Exadata infrastructure, you can also use a backup as the source of your clone. When using a backup, you can select a listed backup to clone from, or create a point-in-time clone. Point-in-time clones contain all data up to a specified timestamp. The specified timestamp must be in the past.

Note

When you create a clone from a backup, you must select a backup that is at least two hours old.

Which statement is NOT correct when using Automatic Indexing on Autonomous Data Warehouse?

- A. Process includes identify candidates, create invisible auto indexes , verify with implement if performance improved and delete unused.
- B. You can use the `dbms_auto_index` package to report on the automatic task and to set your preferences.

- C. It is enabled by default.
- D. You can control behavior at the statement level using the `use_oto_indexes` or `NO_USE_AUTO_INDEXES` hints.
- E. The indexing feature is implemented as an automatic task that runs at a fixed interval.

Correct Answer: C

<https://docs.oracle.com/en/cloud/paas/autonomous-database/adbsa/autonomous-auto-index.html#GUIDD51A4F85-3AE5-4B7D-9B27-32CEDE8DBD79>

Auto Indexing is disabled by default in Autonomous Database

QUESTION 3

Data Guard is enabled for your Autonomous Database and the Lifecycle State field for the primary database indicates that it is stopped.

Which statement is true?

- A. The standby database is terminated.
- B. The standby database is also stopped.
- C. Switchover is automatically initiated.
- D. Failover is automatically initiated.

Correct Answer: B

QUESTION 4

Where can a user's public ssh key be added on the Oracle Cloud Infrastructure Console in order to execute API calls?

- A. On the Autonomous Database Console.
- B. SSH keys are not required in Oracle Cloud Infrastructure.
- C. SSH keys cannot be added from console. They have to be added using REST APIs only.
- D. Navigate to Identity, select Users panel on the console and select "Add Public Key".

Correct Answer: D

-In the Console, click Identity, and then click Users. Locate the user in the list, and then click the user's name to view the details.

-Click Add Public Key.

-Paste the key's value into the window and click Add.

How to Upload the Public Key

You can upload the PEM public key in the Console, which can be accessed by signing in here: <https://cloud.oracle.com>. If you don't have a login and password for the Console, contact an administrator.

1. Open the Console, and sign in.
2. View the details for the user who will be calling the API with the key pair:
 - If you're signed in as the user:
Open the Profile menu (👤) and click User Settings.
 - If you're an administrator doing this for another user: Open the navigation menu and click Identity & Security. Under Identity, click Users. Select the user from the list.
3. Click Add Public Key.
4. Paste the contents of the PEM public key in the dialog box and click Add.

QUESTION 5

An Autonomous Database (ADB) user with an instance wallet has left the company and they shared a database user id with other users when accessing ADB.

Other than changing the shared user password, what can the admin do to protect the instance?

- A. Trust the user who left to not access the database.
- B. Shutdown and restart the instance.
- C. Rotate the instance wallet and share the new wallet with the remaining users.
- D. Delete the user.

Correct Answer: C

Rotate Wallets for Autonomous Database



Wallet rotation lets you invalidate existing client certification keys for a database instance or for all Autonomous Database instances in a region.

You may want to rotate wallets for any of the following reasons:

- If your organization's policies require regular client certification key rotation.
- When a client certification key or a set of keys is suspected to be compromised.

There are two options for client certification key rotation:

- Per-database with **Instance Wallet** selected:
 - For the database whose certification key is rotated, any existing database specific instance wallets will be void. After you rotate a wallet you have to download a new wallet to connect to the database.
 - Regional wallets containing all database certification keys continue to work.
 - All user sessions are terminated for the database whose wallet is rotated. User session termination begins after wallet rotation completes, however this process does not happen immediately.
- Regional level with **Regional Wallet** selected:
 - For the region whose certification key is rotated, both regional and database specific instance wallets will be void. After you rotate a wallet you have to download new regional or instance wallets to connect to any database in the region.
 - All user sessions are terminated for the databases in the region whose wallet is rotated. User session termination begins after wallet rotation completes, however this process does not happen immediately.

QUESTION 6

What command should you issue to enable application continuity on Autonomous Database Shared Infrastructure?

- A. alter database set application_continuity=true
- B. alter session set application_continuity=true
- C. exec dbms_cloud_admin.enable_app_cont (' ');
- D. alter system set application_continuity=true

Correct Answer: C

<https://docs.oracle.com/en/cloud/paas/autonomous-database/adbsa/dbms-cloud-admin.html#GUID42B9DDB0-12A7-4EE3-9860-30EFA41061E9> If it was asking for Autonomous Database Dedicated Infrastructure, then answer would be execute `DBMS_APP_CONT_ADMIN.ENABLE_AC('databaseid_tpurgent.adb.oraclecloud.com', 'LEVEL1', 600);`

QUESTION 7

Which statement is true when you provision an Autonomous Database using the cloning method?

- A. After an Autonomous database is created using a clone, the network access option can be changed.
- B. A clone database source can be a running database instance.
- C. A clone database source can be from a point-in-time backup that is less than 2 hours old.

D. When cloning an Autonomous Database, Oracle Machine Learning workspaces, projects, and notebooks of the source database are cloned to the new database.

Correct Answer: B

QUESTION 8

Which two are correct actions to take in order to Download the Autonomous Database Credentials?

- A. Click on the Autonomous Data Warehouse in the menu, click a database name, then Choose DB Connection button, then Download the Wallet.
- B. Click on the Autonomous Data Warehouse section, pick a database, then Choose Actions, then Download the Wallet.
- C. Click the Compute section of the menu, then choose Instance Configurations, then Download Wallet.
- D. Click on the Object Storage and find your Autonomous Bucket and Download the Wallet Credentials.
- E. Find the Service Console for your Autonomous Database, then pick Administration, then Download the Client Credentials (Wallet).

Correct Answer: AE

<https://docs.oracle.com/en/cloud/paas/autonomous-database/adbsa/connect-download-wallet.html#GUIDB06202D2-0597-41AA-9481-3B174F75D4B1>

QUESTION 9

13 Using Oracle Graph with Autonomous Database



Oracle Graph with Autonomous Database enables you to create graphs from data in your Autonomous Database. With graphs you can analyze your data based on connections and relationships between data entities.

As an Analyst or a Developer you can use graph algorithms and graph pattern queries for ranking, clustering, and path analysis in a graph model of your data. You can use graph features to detect anomalous patterns, identify communities, and find new connections in your data. Then, you can use graphs in your applications, for example, for fraud detection in banking, improved traceability in smart manufacturing, building linked data applications, and more; all while gaining enterprise-grade security, ease of data ingestion, and support for a wide range of workloads.

Autonomous Database includes all the graph capabilities from Oracle Database. In addition, it includes Graph Studio, which further automates graph data management and simplifies modeling, analysis, and visualization across the graph analytics lifecycle.

About Oracle Graph Studio with Autonomous Database

Graph Studio features include automated modeling to create graphs from database tables, an integrated notebook to run graph queries and analytics, and native graph and other visualizations. You can invoke nearly 60 pre-built graph algorithms and visualize your data with many visualization options. Graph Studio is a fully integrated, automated feature with Autonomous Database.

See [Graph Studio: Interactive, Self-Service User Interface](#) and [Access Graph Studio Using Oracle Cloud Infrastructure Console](#) for more information on Graph Studio.



Note: Oracle Autonomous JSON Database does not include Graph Studio.

Which three are use cases for Graph Studio? (Choose three.)

- A. 3-D modelling
- B. Churn analysis
- C. Pattern matching
- D. Facial recognition
- E. Clustering

Correct Answer: ACE

Graph Studio for data scientists

Data scientists need more insights from their data, which can become more accessible through graph analytics and the creation of new engineered features. When it comes to machine learning, data scientists can include those features derived from graphs to generate new insights, such as using **clustering** to find similar customers based on the products they bought.

With Graph Studio, data scientists can efficiently analyze the connectivity in the data and enrich it through feature engineering with the Graph Studio modeler, in-memory graph server, notebooks, and end-to-end analytics flows. Then, data scientists can share the workflow and results collaboratively so others can use it for implementation in a

Through Graph Studio, data scientists can take the graph as an input, adjust it as needed, and execute algorithms and perform pattern-matching queries all within a collaborative notebook environment. Because the graph model does not mandate a fixed schema, the definition of entities and relationships as well as their properties can evolve over time without necessarily impacting all previous work. This allows for faster results and more agile development.



PGQL, a graph query language

Graph Studio provides general-purpose property graph support. PGQL is a powerful SQL-like graph query language. Analysts, developers, and data scientists can also query by using PGQL to search for surrounding nodes, traverse property paths, **pattern matching**, and extracting sub-graphs.

Introducing Graph Studio, part of Oracle Autonomous Database

With the addition of Graph Studio, Oracle Autonomous Database is now a complete, managed platform for analyzing and visualizing graph models.

With Oracle Autonomous Database, you gain a complete graph database platform that can be **deployed in minutes** with one-click provisioning, integrated tooling, and security, which makes graph analytics a possibility even for beginners.



The new comprehensive tooling includes:

- Automated graph modeling
- Extensive graph analytics and graph query support
- Advanced notebooks and integrated visualization
- Automated install, upgrade, and provisioning

Additional new features include:

- Autosave, backup, and checkpoint data restoration features
- Ability to schedule graph analysis
- Sample notebooks and pre-built templates and workflows for different graph use cases

Graph and the converged database

Graph Studio is part of the Autonomous Database, a self-service database and analytics environment that is self-driving, self-securing, and self-repairing.

Because Autonomous Database is a converged database, that means you can seamlessly perform graph analysis on data used in other systems, like data warehouses or transaction systems. You can also transparently use in-memory and partitioning features to enhance query performance and scalability.

QUESTION 10

When integrating object storage with autonomous database how effectively oracle uses the files in the object storage?

- A. Scan partition tables
- B. scan hybrid partition tables
- C. prun column in parquet files
- D. prun columns in csv files

Correct Answer: A

QUESTION 11


When enabling auto-scaling, which three options do not change?

- A. Number of OCPUs displayed in the console.
- B. Amount of storage
- C. I/O resources
- D. parallelism settings
- E. Level of concurrency

Correct Answer: BDE

Explanation: <https://docs.oracle.com/en/cloud/paas/autonomous-database/adbsa/autonomous-autoscale.html#GUID-27FAB1C1-B09F-4A7A-9FB9-5CB8110F7141>

Use Auto Scaling



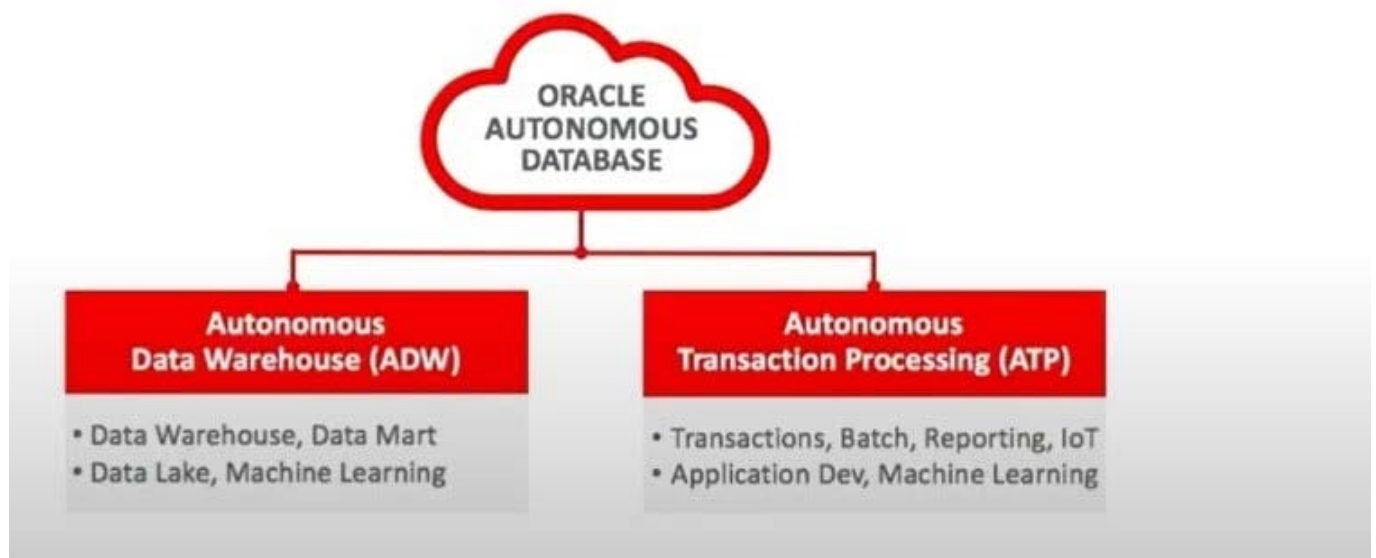
Auto scaling is enabled by default when you create an Autonomous Database instance or you can use **Scale Up/Down** on the Oracle Cloud Infrastructure console to enable or disable auto scaling.

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.

QUESTION 12

AUTONOMOUS DATABASE SERVICES

The underlying converged database capabilities of the Oracle Database enable the Autonomous Database to be offered in two editions that are specifically tailored to a workload following Oracle's Best Practice recommendations. Oracle Autonomous Data Warehouse (ADW) is tailored to Data Warehousing, Data Marts, Data Lakes, and Machine Learning workloads. Oracle Autonomous Transaction Processing (ATP) is tailored to On-Line Transaction Processing, Batch, reporting, IoT, application development, machine learning, and mixed workload environments.



In which four ways can Oracle Database optimally access data in Object Storage?

- A. optimized data archive using partitioned external tables
- B. optimized data archive using hybrid partitioned tables
- C. scan avoidance using block skipping when reading parquet and orc files
- D. scan avoidance using columnar pruning for columnar stores like parquet and orc
- E. scan avoidance using partitioned external tables
- F. scan avoidance using columnar pruning for .csv files

Correct Answer: ABDE