



1Z0-931-22^{Q&As}

Oracle Autonomous Database Cloud 2022 Professional

Pass Oracle 1Z0-931-22 Exam with 100% Guarantee

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

<https://www.pass4itsure.com/1z0-931-22.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**

Which two statements are true about Data Insights?

- A. You can retrieve previously executed Data Insights searches.
- B. Data Insights is available with on-premises deployments of Oracle Database.
- C. Data Insights is an additional priced option.
- D. You can search for Data Insights against a base table or business model.

Correct Answer: AD

QUESTION 2

What is the default port number for accessing a newly created Autonomous Database?

- A. 1524
- B. 1522
- C. 1521
- D. 1523
- E. 1526

Correct Answer: B

<https://www.oracle.com/a/ocom/docs/cloud/autonomous-database-100.pdf>

QUESTION 3

A Corporation is building a web application to allow its customers to schedule service requests online. There is also a need to run operational reports at times during non-peak hours. The architecture team is debating whether such reports should be run on the OLTP database or in a separate data mart. The DBA Manager does not want to add anymore admin responsibility to the team and is looking for a database option that's low to zero maintenance, but meets their strict performance requirements as well.

Which Oracle Cloud Infrastructure database service is appropriate for this scenario?

- A. Since the application needs to be highly available, it should be deployed on a Kubernetes Cluster.
- B. ADW since operational reporting is a higher priority in this scenario.
- C. It is best to build a separate data warehouse, and move the OLTP data on a nightly basis.
- D. ATP. Using 'tpurgent' and 'high' TNS services to separate connection types.

Correct Answer: D



Autonomous Transaction Processing provides all of the performance of the market leading Oracle Database in an environment that is tuned and optimized to meet the demands of a variety of applications, including: mission-critical transaction processing, mixed transactions and analytics, IoT, and JSON document store.

As a service Autonomous Transaction Processing does not require database administration. With Autonomous Transaction Processing you do not need to configure or manage any hardware, or install any software. Autonomous Transaction Processing handles creating the database, backing up the database, patching and upgrading the database, and growing or shrinking the database.

QUESTION 4

Which statement is true when the Autonomous Database has auto scaling enabled?

- A. Database concurrency is scaled up when needed by the workload.
- B. Increases the number of sessions available to the database.
- C. Scales the PGA and SGA size when needed for the workload.
- D. Enables the database to use up to 3x CPU/IO resources immediately when needed by the workload.

Correct Answer: D

QUESTION 5

Which two statements are true when running DBMS_CLOUD.COPY_DATA? (Choose two.)

- A. The source file can be in either Oracle Standard Storage or Oracle Archive Storage bucket in the Object Store.
- B. The source files can reside in Oracle Object Storage, Amazon S3 Object storage, or Azure Blob storage.
- C. A valid credential must be created prior to running the DBMS_CLOUD.COPY_DATA procedure.
- D. The source file will be automatically removed after the DBMS_CLOUD.COPY_DATA procedure finishes successfully.
- E. The target table will be created in Autonomous Database if it does not already exist.

Correct Answer: BC

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/file-uniformats.html#GUID-5D3E1614-ADF2-4DB5-B2B2-D5613F10E4FA>

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/dbms-cloud.html#GUID9428EA51-5DDD-43C2-B1F5-CD348C156122>

QUESTION 6

Which three statements are true about procedures in the DBMS_CLOUD package? (Choose three.)

- A. The DBMS_CLOUD.PUT_OBJECT procedure copies a file from Cloud Object Storage to the Autonomous Data



Warehouse.

B. The DBMS_CLOUD.CREATE_CREDENTIAL procedure stores CloudObject Storage credentials in the Autonomous Data Warehouse database.

C. The DBMS_CLOUD.VALIDATE_EXTERNAL_TABLE procedure validates the source files for an external table, generates log information, and stores the rows that do not match the format options specified for the external table in a badfile table on Autonomous Data Warehouse.

D. The DBMS_CLOUD.DELETE_FILE procedure removes the credentials file from the Autonomous Data Warehouse database.

E. The DBMS_CLOUD.CREATE_EXTERNAL_TABLE procedure creates an external table on files in the cloud. You can run queries on external data from the Autonomous Data Warehouse.

Correct Answer: BCE

DELETE_FILE Procedure

This procedure removes the specified file from the specified directory on Autonomous Data Warehouse.

CREATE_CREDENTIAL Procedure

This procedure stores Cloud Object Storage credentials in the Autonomous Data Warehouse database.

Use stored credentials for data loading or for querying external data residing in the Cloud.

PUT_OBJECT Procedure

This procedure copies a file from Autonomous Data Warehouse to the Cloud Object Storage. The

maximum file size allowed in this procedure is 5 gigabytes (GB).

VALIDATE_EXTERNAL_TABLE Procedure This procedure validates the source files for an external table, generates log information, and stores the rows that do not match the format options specified for the external table in a badfile table on Autonomous Data Warehouse. CREATE_EXTERNAL_TABLE Procedure This procedure creates an external table on files in the Cloud. This allows you to run queries on external data from Autonomous Data Warehouse.

To use Data Pump from ADB, a credential identifying the Object Storage bucket to use must be defined with a DBMS_CLOUD.CREATE_CREDENTIAL function. This will allow ADB to access objects that are stored in the object store, including dump files. To export an existing database to prepare for import into ADB, use the XTP command and add the exclude option for database functionality that is not recommended or supported in ADB. This will prevent errors during the import process. This process is not automatic. And if the logs are not moved, you will receive a warning when running the MDB that the logs are not there. In this example, we're moving the log import.log to object store with a DBMS_CLOUD.PUT_OBJECT command.

VALIDATE_EXTERNAL_TABLE Procedure This procedure validates the source files for an external table, generates log information, and stores the rows that do not match the format options specified for the external table in a badfile table on Autonomous Database. The overloaded form enables you to use the operation_id parameter. PUT_OBJECT Procedure This procedure copies a file from Autonomous Database to the Cloud Object Storage. The maximum file size allowed in this procedure is 5 gigabytes (GB) DELETE_FILE Procedure This procedure removes the specified file from the specified directory on Autonomous Database. CREATE_EXTERNAL_TABLE Procedure This procedure creates an external table on files in the Cloud. This allows you to run queries on external data from Autonomous Database.

QUESTION 7



Oracle Autonomous Database on Dedicated Infrastructure is composed of which Oracle Cloud resources?

- A. Virtual Cloud Network, Compartments, Policies, and Autonomous Exadata infrastructure
- B. Oracle Machine Learning Zeppelin Notebook, Autonomous Exadata infrastructure, Fleet Administrator, and Database Administrator
- C. Autonomous Exadata infrastructure, Autonomous Backup, Autonomous Container Database, and Autonomous Database
- D. Fleet Administrator, Database Administrator, Database User, and Autonomous Exadata infrastructure

Correct Answer: C

QUESTION 8

Which cloud service is NOT supported to import data pump files from using Dedicated Exadata Infrastructure?

- A. Oracle Cloud Infrastructure Object Storage
- B. Google Cloud Storage
- C. Amazon S3
- D. Oracle Cloud Infrastructure Object Storage Classic

Correct Answer: B

QUESTION 9

While Autonomous Transaction Processing and Autonomous Data Warehouse use the same Oracle database, which statement is true about the workloads?

- A. Autonomous Transaction Processing memory usage optimizes workloads for parallel joins and aggregations.
- B. Autonomous Data Warehouse workloads are optimized for mixed workloads.
- C. Autonomous Transaction Processing workloads are optimized for data warehouse, data mart, and data lake.
- D. Data that is bulk loaded, by default, uses the row format in Autonomous Transaction Processing where Autonomous Data Warehouse data format is columnar.

Correct Answer: D

<https://docs.oracle.com/en/cloud/paas/atp-cloud/atpug/autonomous-intro-atp.html#GUID-B90147B2EA04-4147-9606-0413D9CC1589>

<https://docs.oracle.com/en/cloud/paas/atp-cloud/atpug/autonomous-intro-atp.html#GUID-B90147B2EA04-4147-9606-0413D9CC1589>

QUESTION 10



Which statement is FALSE about setting up compartment quotas?

- A. Quotas set on a parent compartment override quotas set on child compartments.
- B. Compartment quotas cannot be set on the root compartment.
- C. Compartment quotas use policies that allow allocation of resources with a high level of flexibility.
- D. You cannot manage quotas in a compartment, unless you belong to a group that has the correct permissions.

Correct Answer: B

QUESTION 11

Which can be Scaled Independently of the number of CPUs In an Autonomous Database?

- A. Parallelism
- B. Storage
- C. Concurrency
- D. Memory
- E. Sessions

Correct Answer: B

<https://docs.oracle.com/en/cloud/paas/autonomous-database/adbsa/autonomous-keyfeatures.html#GUID-7D0627D9-1B40-40B1-865A-84FF65A65846> Key Features Fully elastic scaling: Scale compute and storage independently to fit your database workload with no downtime:

- **Fully elastic scaling:** Scale compute and storage independently to fit your database workload with no downtime:
 - Size the Autonomous Database to the exact compute and storage required
 - Scale the Autonomous Database on demand: Independently scale compute or storage
 - Shut off idle compute to save money
-

QUESTION 12

Your customer has upgraded their on premise 11.2 database to 12.2. During this migration the database was migrated to a pluggable database and is now in production. How should the customer unplug their database to migrate to Autonomous Database (ADB)?

- A. Pluggable databases cannot be migrated to ADB using plug, unplug or clone.
- B. Create a database link from source database to the ADB environment and clone the PDB.
- C. Unplug to an xml file so database files and xml file can be uploaded to object storage.
- D. Unplug into a pdb archive which can be uploaded to object storage.



Correct Answer: A

Since an ADB database has some restrictions on the object types and Oracle Database Options you need to use a logical migration method rather than a physical one.

The main migration tool for migrating to ADB is Data Pump. You can export your schemas and import them into ADB using Data Pump. To sync up the additional/incremental changes on the source database during the export/import process you can use GoldenGate or GoldenGate Cloud Service to replicate those changes to ADB.

In the current release you cannot use physical migration methods like backup/restore, Data Guard, Pluggable/Unplug, database clones, and transportable tablespaces to move your existing database to ADB.

QUESTION 13

Autonomous Database's auto scaling feature allows your database to use up to three times the current base number of CPU cores at any time. As demand increases, auto scaling automatically increases the number of cores in use. Likewise, as demand drops, auto scaling automatically decreases the number of cores in use.

Which statement is FALSE regarding the auto scaling feature?

- A. For databases on dedicated Exadata infrastructure, the maximum number of cores is limited by the number of free cores that aren't being used by other auto scaling databases to meet high-load demands.
- B. The base number of OCPU cores allocated to a database is not guaranteed.
- C. For databases on dedicated Exadata infrastructure, the maximum number of cores available to a database depends on the total number of cores available in the Exadata infrastructure instance.
- D. Auto scaling is enabled by default and can be enabled or disabled at any time.


Correct Answer: B

QUESTION 14



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

QUESTION 15

Which three of the following data sources are available when using the Data Load page on Database Actions?

- A. Files in Oracle Cloud Storage
- B. Backup files in block storage
- C. Local Files
- D. Files in AWS S3 Storage



E. REST endpoints

Correct Answer: ACD

[1Z0-931-22 VCE Dumps](#)

[1Z0-931-22 Practice Test](#)

[1Z0-931-22 Study Guide](#)