



Oracle

1Z0-931

Oracle Autonomous Database Cloud Certified Specialist(1z0-931)

Version 1
Updated: September 03, 2019

[Total Questions: 50+]

Web: www.k21academy.com

Email: support@k21academy.com

Blog: <http://k21academy.com/blog>



Atul Kumar

Oracle ACE, Author & Cloud Expert

Module 1: OCI Database Overview

M01Q01 Where can a user's public ssh key be added on the Oracle Cloud Infrastructure Console in order to execute API calls?

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

Answer: A

Explanation:

The screenshot displays the OCI console interface for a user named 'K21AcademyUser'. The user's profile includes a green circular avatar with a white 'K', the status 'ACTIVE', and a description 'OCI New User'. Below this, there are tabs for 'User Information' and 'Tags'. The 'User Information' tab shows details such as OCID, creation time, multi-factor authentication status, and email. The 'Capabilities' section lists various permissions like local password, API keys, and auth tokens. At the bottom, the 'API Keys' section indicates 'No API Keys' and features a blue 'Add Public Key' button, which is highlighted with a red box in the original image.

M01Q02 Which two statements are true about the The Oracle Cloud Infrastructure (OCI)?

- A. Because availability domains do not share infrastructure such as power or cooling, or the internal availability domain network, a failure at one availability domain within a region is unlikely to impact the availability of the others within the same region.
- B. Regions are dependent on other regions and must be located within 5 thousand kilometres of each other.
- C. An OCI region is a localized geographic area, and an availability domain is one or more data centres located within a region.
- D. A single fault domain can be associated with multiple regions and availability domains.

support@k21academy.com

<https://k21academy.com/autonomous03>

Answer: A, C

Explanation:

A is correct as Availability domains are isolated from each other, fault tolerant, and very unlikely to fail simultaneously. Because availability domains do not share infrastructure such as power or cooling, or the internal availability domain network, a failure at one availability domain within a region is unlikely to impact the availability of the others within the same region.

B is wrong as Regions are independent of other regions and can be separated by vast distances—across countries or even continents

C is correct as A region is a localized geographic area, and an availability domain is one or more data centres located within a region

D is wrong as fault domain is a grouping of hardware and infrastructure within an availability domain

Reference: <https://docs.cloud.oracle.com/iaas/Content/General/Concepts/regions.htm>

Module 2: Autonomous Database Overview

M02Q01 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. ATP. Using 'tpurgent and 'high' TNS services to separate connection types
- B. It is best to build a separate data warehouse, and move the OLTP data on a nightly basis.
- C. Since the application needs to be highly available, It should be deployed on a Kubernetes Cluster.
- D. ADW since operational reporting is a higher priority in this scenario.

Answer: A

Explanation:

Since there is a requirement of zero admin responsibility and strict performance and also run reports on top of OLTP database, ATP is the best suited for this which has tpurgent & high tns services to restrict performance on requirement.

M02Q03 Which three statements are true regarding how Autonomous Database provides data security?

- A. Data is encrypted at rest using transparent data encryption.
- B. Oracle automatically applies security updates to ensure data is not vulnerable to known attack vectors.
- C. Users are given OS logins or SYSDBA privileges to prevent phishing attacks.
- D. Network connections from clients to Autonomous Database are encrypted using the client credentials wallet.

Answer: A, B, D

Explanation:

A is correct as Data is encrypted using transparent data encryption

B is correct as Autonomous comes with automated patching, so every patches are applied and taken care of by Oracle

C is wrong as in Autonomous Database we don't have access to OS or SYSDB

D is correct as in order to connect to autonomous database you need to use credentials wallet

Reference: <https://docs.oracle.com/en/cloud/paas/atp-cloud/atpug/autonomous-security.html#GUID-B26111FC-311F-4A80-96F7-E3608BEC25C7>

M02Q04 Which is correct about security features that are available in Oracle Autonomous Database?

- A. TDE but not Data Redaction
- B. Neither Data Redaction nor TDE are supported .
- C. Data Redaction and TDE are both supported
- D. Data Redaction but not TDE

Answer: C

Explanation:

A, B & D is wrong TDE & both Data Redaction are supported in Autonomous Database

Reference: <https://www.oracle.com/database/technologies/datawarehouse-bigdata/adbfags.html>

M02Q05 Which task is NOT automatically performed by the Oracle Autonomous Database?

- A. Mask your sensitive data
- B. Patching the database
- C. Backing up the database
- D. Automatically optimize the workload

Answer: D

Explanation:

A is correct as masking sensitive data is taken care of by you as a customer

B,C & D is wrong as these are automated and taken care of by Autonomous Database

Reference: <https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/autonomous-key-features.html#GUID-7D0627D9-1B40-40B1-865A-84FF65A65846>

M02Q06 Which three tasks by default are taken care of by the Autonomous Database?

- A. Firmware Patching
- B. Database Upgrades
- C. Backups
- D. Application User Creation
- E. Data Loading

Answer: A, B, C

Explanation:

A, B & C are correct as these tasks are taken care of by Autonomous Database
D & E are wrong as it is your duty to create users or load data to Autonomous Database

Reference: <https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/autonomous-key-features.html#GUID-7D0627D9-1B40-40B1-865A-84FF65A65846>

M02Q07 Autonomous Data Warehouse configures and optimizes your database for you, therefore you do not need to perform some of the administration operations for configuring the database. Which statement is true regarding the default configuration of the Autonomous Data Warehouse?

- A. In Autonomous Data Warehouse direct access to the database node and the local file system is allowed.
- B. Degree of parallelism for SQL statements is set based on the number of OCPUs in the system and the database service the user is connecting to.
- C. Compression is not enabled by default and you must specify the compression method for your tables with the compression clause in your CREATE TABLE and ALTER TABLE commands.
- D. Adding, removing, or modifying tablespaces is allowed.

Answer: B

Explanation:

A is wrong as in ADW you don't have any access to the DB node and the local file system, you just connect connect to the Autonomous database using Database client using ADMIN or other user which has permission to connect to ADW, as you also don't have sys dba privilege.

B is correct as Parallelism is enabled by default. Degree of parallelism for SQL statements is set based on the number of OCPUs in the system and the database service the user is connecting to.

C is not correct as Compression is enabled by default, Autonomous Data Warehouse uses Hybrid Columnar Compression for all tables by default.

D is wrong as The default data and temporary tablespaces for the database are configured automatically. Adding, removing, or modifying tablespaces is not allowed.

Reference: <https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/experienced-database-users.html#GUID-58EE6599-6DB4-4F8E-816D-0422377857E5>

M02Q08 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 ie. data format is columnar.

Answer: D

Explanation:

A is wrong as in ATP memory usage optimizes for critical working set can be cached to avoid IO And in case of ADW memory usage optimizes workloads for parallel joins and aggregations.
B is incorrect as Autonomous Data Warehouse is tuned and optimized for data warehouse workloads.
C is wrong as Autonomous Transaction Processing is tuned and optimized for transaction processing workloads.
D is correct as by default in ATP data format is row and in ADW data format is Columnar

Reference: <https://docs.oracle.com/en/cloud/paas/atp-cloud/atpug/autonomous-intro-atp.html#GUID-B90147B2-EA04-4147-9606-0413D9CC1589>

<https://docs.oracle.com/en/cloud/paas/atp-cloud/atpug/autonomous-intro-atp.html#GUID-B90147B2-EA04-4147-9606-0413D9CC1589>

M02Q09 What are two security features enabled by default by the Autonomous Database?

- A. One SYSDBA account
- B. Transparent Data Encryption
- C. SQL Net Encryption
- D. Encrypted Database Links
- E. Read Only access to OS Audit logs

Answer: B, C

Explanation:

A is wrong as there is no SYSDBA privilege to user

B is correct as in ADB All data encrypted at rest using transparent data encryption.

C is correct as Autonomous Database uses Oracle Net Services (also known as SQL*Net), which uses an encrypted key stored in a *wallet* on both the client (where the application is running) and the server (where your database service on the Autonomous Data Warehouse is running).

D is wrong as there is no encrypted DB Links in Autonomous Database wherein To ensure security, the database link port is restricted to the range 1521–1525

Reference:

<https://www.oracle.com/database/technologies/datawarehouse-bigdata/adb-faqs.html#SECURITY-BOOKMARK>

<https://docs.oracle.com/en/cloud/paas/atp-cloud/atpug/database-links.html#GUID-28F06977-300F-4599-AD97-2DEF67EF5FEA>

M02Q10 Which of these database features is NOT part of the Autonomous Database?

- A. Flashback Database
- B. Online Indexing
- C. Real Application Clusters (RAC)
- D. Java in the Database

Answer: D

Explanation:

A, B & C are incorrect because Flashback Database, Online Indexing & RAC are part of Autonomous Database

D is correct as Java in DB is not the part of Autonomous Database

Reference:

<https://docs.oracle.com/en/cloud/paas/atp-cloud/atpuq/experienced-database-users.html#GUID-B6FB5EFC-4828-43F4-BA63-72DA74FFDB87>

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/experienced-database-users.html#GUID-5CF411BF-69E3-43C1-9379-25257A88A354>

Module 3: Provisioning & Connecting Autonomous Database

M03Q01 Which Java Database Connectivity (JDBC) thin client version supports the use of HTTP proxy settings that does NOT support TCP to connect to Autonomous Data Warehouse?

- A. 12c Release 2 onwards
- B. Any version
- C. 18.1 onwards
- D. 17.4 onwards

Answer: A

Explanation:

Connections with an HTTP Proxy

If the client is behind a firewall and your network configuration requires an HTTP proxy to connect to the internet, then perform the following steps to update the `sqlnet.ora` and `tnsnames.ora` files.

Note:

Connections through an HTTP proxy are only available with Oracle Client software version 12.2.0.1 or later.

Reference:

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/connect-preparing.html#GUID-EFAFA00E-54CC-47C7-8C71-E7868279EF3B>

M03Q02 Which Autonomous Database Service is NOT used to connect to an Autonomous Transaction Processing instance?

- A. TPURGENT
- B. TPPERFORMANT
- C. MEDIUM
- D. LOW
- E. HIGH

Answer: B

Explanation:

A, C, D, E are wrong as all these database services is used to Connect to a Transaction Processing instance

B is correct as there is no Autonomous Database service named as TPPERFORMANT

Reference:


<https://docs.oracle.com/en/cloud/paas/atp-cloud/atpug/manage-priorities.html#GUID-80E464A7-8ED4-45BB-A7D6-E201DD4107B7>

M03Q03 Which is NOT required to connect to Autonomous Database from SQL developer?

- A. Service name
- B. Username and password
- C. Wallet file
- D. Database name

Answer: D

Explanation:



Reference:

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/connect-sql-dev182.html#GUID-14217939-3E8F-4782-BFF2-021199A908FD>

M03Q04 If you need to connect to Autonomous Data Warehouse (ADW) using Java Database Connectivity (JDBC) via an HTTP proxy, where do you set the proxy details?

- A. tnsnames.ora
- B. cwallet.sso
- C. ojdbc.properties
- D. keystore.jks
- E. sqlnet.ora

Answer: A

Explanation:

A is correct as we add all the details such as HTTP proxy hostname and port to the connection definitions in tnsnames.ora.

B, C, & D are not required to enable HTTP proxy

E is wrong as we just enable HTTP proxy details from it and not add the proxy details

Reference:

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/connect-jdbc-thin-wallet.html#GUID-D583E056-CB47-4860-B303-039DEDEC49B8>

M03Q05 Which statement is true regarding database client credentials file required to connect to your Autonomous Database?

- A. The Transparent Data Encryption (TDE) wallet can be used for your client credentials to connect to your database.
- B. Store credential files in a secure location and share the files only with authorized users to prevent unauthorized access to the database.
- C. When you share the credential files with authorized users, mail the wallet password and the file in the same email.
- D. Place the credential files on a share drive that all users can use to connect to the database.

Answer: B

Explanation:

A is wrong TDE wallet is not used to connect to Autonomous Database

B is correct as Store wallet files in a secure location. Share wallet files only with authorized users. If wallet files are transmitted in a way that might be accessed by unauthorized users (for example, over public email), transmit the wallet password separately and securely.
C is wrong as we only required password while downloading the wallet, once it is downloaded it can be shared and accessible to authorized users without any requirement of password
D is wrong as we cannot share it with everyone apart from authorized user as that will be security concern

Reference:

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/connect-download-wallet.html#GUID-B06202D2-0597-41AA-9481-3B174F75D4B1>

M03Q06 How many predefined service names are configured in tnsnames.ora for a single Autonomous Transaction Processing database instance, and what are they called?

- A. None. There are no predefined service names in tnsnames.ora.
- B. Three. They are called high, medium and low.
- C. Two. They are called ATP and ADW.
- D. Five. They are called tpurgent, tp, high, medium and low.

Answer: D

Explanation:

What pre-defined database services are available with ATP?

- ATP comes with 5 preconfigured database services, HIGH, MEDIUM, LOW, TP and TPURGENT. The services control the priority of the sessions when the system is under resource pressure and some control the parallel degree used. By default, queries will execute serially when connected to the TP, TPURGENT and LOW services.

Reference:

<https://www.oracle.com/database/technologies/datawarehouse-bigdata/adw-faqs.html#ATP>

Module 4: Administer Autonomous Database (Start | Stop| Scale | Terminate | Backup Restore| Clone)

M04Q01 When you choose to scale your Autonomous Database, which statement is true in regards to OCPUs and storage?

- A. OCPUs and storage can be scaled independently.
- B. Storage in TB cannot be larger than the number of OCPUs.
- C. OCPUs and storage must remain in sync.
- D. Increasing OCPUs will automatically increase storage.

Answer: A

Explanation:

A is correct as both CPU & Storage can be scale -up independently
B is wrong as it is not necessary that Storage in TB should be greater than number of OCPUs
C is wrong as it does not required for Storage & OPCU to be in sync
D Is wrong as increasing OCPUs will not automatically increase storage.

Reference:

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/autonomous-add-resources.html#GUID-DA72422A-5A70-42FA-A363-AB269600D4B0>

M04Q03 Which can be Scaled independently of the number of CPUs in an Autonomous Database?

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

Answer: D

Explanation:

A, B, C & E is wrong as all these are dependent on CPUs in an Autonomous Database
D is correct as you can independently scale the number of CPUs in Autonomous Database

Reference:

<https://www.oracle.com/database/technologies/datawarehouse-bigdata/adb-faqs.html#GENERAL>

M04Q04 What two tasks can be executed from the service console for Autonomous Databases?

- A. Creating and scaling of Autonomous Database service.
- B. Creating schemas.
- C. Autonomous Databases monitoring for usage and query performance.
- D. Wizard to download connection wallet for connection from desktop tools.

Answer: C, D

Explanation:

Autonomous Data Warehouse

Overview

Activity

Administration

Development

DATABASE

DB201908281558

Storage used

0%

2.59 GB / 1 TB

CPU utilization (%)

No data to display

Running SQL statements

No data to display

Number of OCPUs allocated

No data to display

SQL statement response time (s)

No data to display

Autonomous Data Warehouse

Overview

Activity

Administration

Development

DATABASE

DB201908281558

Download Client Credentials (Wallet)

Connections to Autonomous Data Warehouse use a secure connection. Your existing tools and applications will need to use this wallet file to connect to your Autonomous Data Warehouse Instance. If you are familiar with using an Oracle Database within your own data center, you may not have previously used these secure connections.

Set Resource Management Rules

Set resource management rules to allocate CPU/IO shares to consumer groups and to cancel SQL statements based on their runtime and amount of IO.

Set Administrator Password

Set or reset your database administrator user's (ADMIN) password and when locked unlock your administrator user account on Autonomous Data Warehouse.

Manage Oracle ML Users

Create new Oracle Machine Learning user accounts and manage the credentials for existing Oracle Machine Learning users.

Send Feedback to Oracle

Use our CloudCustomerConnect forum to provide feedback about the service to Oracle, post questions, connect with experts, and share your thoughts and ideas. Click here to link to the forum.

M04Q05 When scaling OCPUs in Autonomous Database, which statement is true in regards to active transactions?

- A. Active transactions are terminated and rolled back.
- B. Active transactions are paused.
- C. Scaling cannot happen while there are active transactions in the database.
- D. Active transactions continue running unaffected.

Answer: D

Explanation:

Autonomous Database: Starting, Stopping, Scaling


- Scale your Autonomous Database Instance on demand without tedious manual steps
 - Independently scale compute or storage
 - Resizing occurs instantly, fully online
 - Memory, IO bandwidth, concurrency scales linearly with CPU
 - **Active transactions continue running unaffected**

M04Q06 Which two options are available to restore an Autonomous Data Warehouse?

- A. Select the snapshot of the backup.
- B. Select the archived redo logs.
- C. Specify the point in time (timestamp) to restore
- D. Select the backup from which restore needs to be done.
- E. Backup and recovery must be done using Recovery Manager(RMAN).

Answer: C, D

Explanation:



References:

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/backup-recover.html#GUID-78C28C41-AA87-4AD7-BEB6-693235C73F3C>

Module 5: Data Loading In Autonomous Database


M05Q01 Your customer receives information in various formats like .csv files from their suppliers. The business user would like to collect all of this information and store it in an ATP environment. The Oracle adviser recommends to use Oracle Data Sync for this. Which statement is true regarding Oracle Data Sync?

- A. Data Sync can only load files into tables (insert-only), the customer has to write the additional code.
- B. Data Sync cannot transform your data while loading it into the destination table.
- C. Data Sync can load a combination of data source, such as xsv, .xlsx and Oracle relational files.
- D. Data Sync can only load data from one source into one destination table.

Answer: C

Explanation:

A is wrong as you can use Data Sync to load file into tables as well as into data sets
B is wrong as Data Sync Can replicate and transform data when the target ADB
C is correct as You can load data from files (CSV and XLSX), various relational sources (tables, views, SQL statements), OTBI, JDBC data sources, and Oracle Service Cloud
D is wrong as you can load data from multiple sources such as MongoDB, Sybase, Hive SalesForce.com, AWS (Redshift and Postgres), MySQL etc to multiple target such as ADB, Essbase CS, BICS/OAC Data Sets etc



M05Q02 Which two statements are true when running DBMS_CLOUD.COPY_DATA?

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

Answer: A, E

Explanation:

- A is correct as valid credential must be created prior to running the DBMS_CLOUD.COPY_DATA procedure
- B is wrong as source file can be reside Oracle Object Storage, Amazon S3 Object storage, or Azure Blob storage.
- C is wrong as you need to define the table name while loading data using DBMS_CLOUD.COPY_DATA
- D is wrong as source file will not be automatically removed after the DBMS_CLOUD.COPY_DATA procedure finishes successfully.
- E is correct as source files can reside in Oracle Object Storage, Amazon S3 Object storage, or Azure Blob storage.

Reference: <https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/dbms-cloud.html#GUID-9428EA51-5DDD-43C2-B1F5-CD348C156122>

M05Q03 In which way can a SQL Developer help you test your data loading scenario to Autonomous Database (ADB)?

- A. In the TEST phase, a temporary table will be populated with the records before inserting them in the destination table.
- B. In the TEST phase of the wizard, a subset of accepted records are displayed based on your definition.
- C. In the TEST phase of the wizard a list is generated containing the records that would be rejected during import.
- D. In the Column Definition Phase, the system cross-references with the file-contents and shows the conflicts with the definition.

Answer: B

Explanation:

SQL Developer and Data Migration

- Use the wizard to export/move/import into Autonomous Database
- Supports many file types such as .dmp and .csv
- Can import the files located from many locations
 - The local system where SQL Developer is located
 - Oracle Object Store
 - Amazon S3 Object Store
- A TEST option of the Wizard can display a subset of records to be imported based on your definition

M05Q04 Which two statements are true with regards to Oracle Data Sync?

- A. Data Sync can use a normal OCI (thick) client connection to connect to an Oracle database.
- B. Data Sync has default drivers available that supported loading data from DB2, Microsoft SQL Server, MySQL and Teradata.
- C. Data Sync can connect to any jdbc compatible source like MongoDB, RedShift and Sybase.
- D. Data Sync can load your data in parallel in order to speed up the loading process.

Answer: C, D

Explanation:

A is wrong as Data Sync uses Java DB Repository to connect to an Oracle database

B is wrong as you need to copy the JDBC driver to for each vendor such DB2, Microsoft SQL Server, MySQL & Teradata to load data from it.

C is correct as Data Sync can connect to any jdbc compatible source like MongoDB, RedShift and Sybase

D is correct as Data Sync Efficient writes parallelizing data loads and using parallel streams for a target table

Reference:

<http://www.ateam-oracle.com/configuring-the-data-sync-tool-for-bi-cloud-service-bics>

M05Q05 What are three methods to load data into the Autonomous Database?

- A. SQL*Loader
- B. Transportable Tablespace
- C. Oracle Data Pump
- D. RMAN Restore
- E. Oracle GoldenGate

Answer: A, C, E

Explanation:

B & D are wrong as transportal tablespace & Oracle Data Pump is not supported in Autonomous database

What are the options for loading data into Oracle Autonomous Database?

- There are multiple ways to load data into ADB:

ADB is integrated with multiple object storage services for data loading. You can upload your source files to one of these object-stores and use the PL/SQL API DBMS_CLOUD to load data into your database. ADB supports Oracle Cloud Infrastructure Object Storage, Oracle Cloud Infrastructure Object Storage Classic, AWS S3, and Azure Blob Storage for loading data. This is the recommended method for loading large data sets. SQL Developer also provides a data loading wizard that can load data from these object stores. DBMS_CLOUD to load data is not applicable for ATP-D because DBMS_CLOUD is not available on ATP-D.

You can use Data Pump Import to load data from dump files. Data Pump Import is integrated with Oracle Cloud Infrastructure Object Storage, Oracle Cloud Infrastructure Object Storage Classic, AWS S3, Microsoft Azure Blob Storage. You can upload your dump files to one of these object-stores and import from those dump files.

You can also use client-side tools like SQL*Loader, or your in-house scripts to load data from files residing on your client machines.

Another option is to use Oracle tools and cloud services like GoldenGate, ODI, Goldengate Cloud Service, and Data Integration Platform Cloud.

Any 3rd party ETL tool that can connect to ADB using the security credentials wallet can also be used to load data.

Reference: <https://www.oracle.com/database/technologies/datawarehouse-bigdata/adbfags.html#ATPD>

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/load-data-intro.html#GUID-5D2F70D8-4FA1-482C-BFB0-43441FB897F3>

Module 6: Managing Users In Autonomous Database

M06Q01 What predefined user is created when an Autonomous Database (ADB) instance is created that you connect to in order to create other users and grant roles?

- A. DWDEV
- B. ADMIN
- C. SYS
- D. SCOTT

Answer: B

Explanation:

ADMIN is the only user which is created when an Autonomous Database (ADB) instance is created and which has all other privileges to manage the Autonomous Database

M06Q02 Which two methods can you use to create users and grant roles in Autonomous Database services?

- A. through SQLPlus
- B. through SQL Developer
- C. using DBMS_CLOUD_ADMIN package
- D. through the Oracle Cloud Infrastructure service console

Answer: A, B

Explanation:

A & B is correct as you can to Autonomous database either using SQLPlus or SQL Developer using admin users and able to create users & grant roles in Autonomous Database Services. C & D is wrong as DBMS_Cloud_ADMIN package is not used to create users and grant roles & using Oracle Cloud Infrastructure Service Console you can create users for OCI not for Autonomous Database

Reference: <https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/connect-sqlplus.html#GUID-A3005A6E-9ECF-40CB-8EFC-D1CFF664EC5A>

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/connect-sql-dev.html#GUID-FB611317-2999-4C8C-8A3B-B3B3D3C5B417>

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/dbms-cloud-admin.html#GUID-D76B229E-781E-45C0-9F14-CAF30F9E6E3B>

Module 7: Machine Learning In Autonomous Database

M07Q01 When exporting a notebook, what type of file is created?

- A. XML
- B. SQL
- C. ASCII
- D. JSON
- E. TXT

Answer: D

Explanation:

About Notebooks

The Notebooks page lists all the notebooks associated with the selected project. You can create, edit, and run your notebooks here.

You can perform the following tasks in the Notebooks page:

- To create a new notebook, select a project and click **Create**.
- To edit a notebook, select the notebook and click **Edit**. You can edit the notebook name, and add comments in the Edit Notebook dialog box.
- To create a copy of a notebook, select the notebook and click **Duplicate**. The duplicate copy of the selected notebook is created, and listed in the Notebooks page with the suffix _1 in the notebook name.
- To save a notebook as a template, select the notebook and click **Save as Template**. You can save the template in **Personal** or **Shared** under Templates.
- To delete a notebook, select it and click **Delete**.
- To import a notebook as **.json files**, click **Import**. Select the project and workspace in which to import the notebook.
- To create versions of a notebook, select it and click **Version**. You can experiment with your notebook by creating versions of it, and revert to an older version by clicking **Revert Version**.
- To open a notebook and run it, click the notebook. The notebook opens in the edit mode.

Reference:

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/omlug/get-started-notebooks-data-analysis-and-data-visualization.html#GUID-009B9073-4B45-4505-8DAD-F1F04FD3E160>

M07Q02 Given the steps:

1. Create Oracle Machine Learning User
2. Create projects
3. Create workspaces
4. Create Notebooks
5. Run SQL Scripts

Which two steps are out of order when working with Oracle Machine learning?

- A. Run SQL Scripts

- B. Create Notebooks
- C. Create projects
- D. Create workspaces
- E. Create Oracle Machine Learning User

Answer: C, D

Explanation:

Workflow for using Oracle Machine Learning

Tasks
1. Create Oracle Machine Learning user
2. Create workspaces
3. Create projects
4. Create notebooks
5. Run SQL scripts/statements

Reference:

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/omlug/get-started-oracle-machine-learning.html#GUID-13452045-6AF6-402B-83CC-C9AFBCEC30CA>

M07Q03 What two methods can you use to define Machine Learning Users?

- A. Client tools
- B. Oracle Cloud Infrastructure Console
- C. SQL/Developer
- D. Use DBMS_CLOUD_ADMIN package

Answer: A, C

Explanation:

Initially, the Role field shows the role None for existing database users. After adding a user the role Developer is assigned to the user.
Use SQL Dev or Client tools

Reference:

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/manage-users-oml.html#GUID-D928B760-513A-4C6D-9EB7-DF992249A26D>

Module 8: Monitoring & Performance

Autonomous Database

M08Q01 Which Autonomous Database Cloud service ignores hints in SQL Statements by default?

- A. Autonomous Data Warehouse
- B. Both services ignore hints by default
- C. Neither service ignores hints by default
- D. Autonomous Transaction Processing

Answer: B

Explanation:

Hints ignored by default by both ADW & ATP, wherein Users can enable hints explicitly

References:

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/manage-optimizer-stats.html#GUID-69906542-4DF6-4759-ABC1-1817D77BDB02>

<https://docs.oracle.com/en/cloud/paas/atp-cloud/atpug/manage-optimizer-stats.html#GUID-69906542-4DF6-4759-ABC1-1817D77BDB02>

M08Q02 When running lots of high concurrent queries in Autonomous Database, you should use which predefined service connection?

- A. DBNAME_CONCURRENT
- B. DBNAME_LOW
- C. DBNAME_HIGH
- D. DBNAME_MEDIUM

Answer: D

Explanation:

A is wrong as there is no predefined service named as "DBNAME_CONCURRENT"

support@k21academy.com

<https://k21academy.com/autonomous03>

B is correct as we need to have high concurrency and low is having highest concurrency
C is wrong as high has lowest concurrency
D is wrong as medium has high concurrency wherein low has the highest one

Reference:

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/manage-priorities.html#GUID-80E464A7-8ED4-45BB-A7D6-E201DD4107B7>

<https://docs.oracle.com/en/cloud/paas/atp-cloud/atpug/manage-priorities.html#GUID-80E464A7-8ED4-45BB-A7D6-E201DD4107B7>

M08Q03 The default eight-day retention period for Autonomous Database performance data can be modified using which DBMS_WORKLOAD_REPOSITORY subprogram procedure?

- A. UPDATE_OBJECT_INFO
- B. MODIFY_SNAPSHOT_SETTINGS
- C. UPDATE_OBJECT_INFO
- D. CREATE_BASELINE_TEMPLATE

Answer: B

Explanation:

The retention time can be changed by modifying the Automatic Workload Repository retention setting with the PL/SQL procedure DBMS_WORKLOAD_REPOSITORY.MODIFY_SNAPSHOT_SETTINGS. Be aware that increasing the retention time will result in more storage usage for performance data

Reference:

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/monitor-performance-intro.html#GUID-4A1E8FCA-FAF1-43F5-93E0-559C0145D6C3>

M08Q04 As a database architect you are tasked with configuring a high concurrency, production OLTP application to connect to an Autonomous Transaction Processing database with a requirement to have some reporting queries run in parallel mode. Which connection service is appropriate for such a workload?

- A. TP
- B. TPURGENT

- C. HIGH
D. MEDIUM

Answer: D

Explanation:

A is wrong as TP does not run with parallelism.

B is wrong as TPURGENT connection service is used for time critical transaction processing operations & its support manual parallelism wherein we have requirement of connection service to support reporting purposes

C is wrong as High provides less concurrency i.e. 3 QUERIES wherein medium provides 1.25 X CPU_COUNT QUERIES

D is correct as Medium is A typical application connection service for reporting and batch operations. All operations run in parallel and are subject to queuing. Using this service the degree of parallelism is limited to four (4) and concurrency is 1.25 X CPU_COUNT QUERIES

SERVICES NAME	LOCATION	PARALELLISM	RESOURCE MANAGEMENT PLAN SHARES	CONCURRENCY
TPURGENT	ATP ONLY	MANUAL	12	UNLIMITED
TP	ATP ONLY	1	8	UNLIMITED
HIGH	ADW & ATP	CPU_COUNT	4	3 QUERIES
MEDIUM	ADW & ATP	4	2	1.25 X CPU_COUNT QUERIES
LOW	ADW & ATP	1	1	2 X CPU_COUNT QUERIES

Reference:

<https://docs.oracle.com/en/cloud/paas/atp-cloud/atpug/manage-priorities.html#GUID-80E464A7-8ED4-45BB-A7D6-E201DD4107B7>

M08Q05 Users are required to select a service when connecting to Autonomous Data Warehouse and these services match to one of three different consumer groups: High, Medium, and Low. Which statement about these consumer groups is correct?

- A. High provides highest resources, lowest concurrency, and DoP is 1.
- B. Low provides highest concurrency, lowest resources, and DoP =1.
- C. Medium provides intermediate resource and concurrency, and queries run in a serial.
- D. High provides highest concurrency and lowest resources, and DoP is i.

Answer: B

Explanation:

A is wrong as Highest resources, lowest concurrency and parallelism depends on CPU Count
B is correct as Low provides highest concurrency, lowest resources, and DoP =1.
C is wrong as Medium provides less resources, higher concurrency. Queries run in parallel.

SERVICES NAME	LOCATION	PARALELLISM	RESOURCE MANAGEMENT PLAN SHARES	CONCURRENCY
TPURGENT	ATP ONLY	MANUAL	12	UNLIMITED
TP	ATP ONLY	1	8	UNLIMITED
HIGH	ADW & ATP	CPU_COUNT	4	3 QUERIES
MEDIUM	ADW & ATP	4	2	1.25 X CPU_COUNT QUERIES
LOW	ADW & ATP	1	1	2 X CPU_COUNT QUERIES

Reference:

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/manage-priorities.html#GUID-80E464A7-8ED4-45BB-A7D6-E201DD4107B7>

Module 9: Managing Autonomous Database Using REST API | CLI | Terraform

M09Q01 What REST verb is used to create an Autonomous Database service using REST APIs?

- A. A "PUT" REST call
- B. A "GET" REST call
- C. A "POST" REST call
- D. An "INSERT" REST call

Answer: C

Explanation:

Example I - creates a new Autonomous Database:

```
POST /20160918/autonomousDatabases
Host: database.us-phoenix-1.oraclecloud.com
<authorization and other headers>
{
  "compartmentId" : "ocid.compartment.oc1..<unique_ID>",
  "displayName" : "example_autonomous_database",
  "dbName" : "adatabasedb1",
  "adminPassword" : "<password>",
  "cpuCoreCount" : 8,
  "dataStorageSizeInTBs" : 1
}
```

Example II - creates an Autonomous Database by cloning an existing one:

```
POST /20160918/autonomousDatabases
Host: database.us-phoenix-1.oraclecloud.com
<authorization and other headers>
{
  "compartmentId" : "ocid.compartment.oc1..<unique_ID>",
  "displayName" : "example_autonomous_database",
  "dbName" : "adatabasedb1",
  "adminPassword" : "<password>",
  "cpuCoreCount" : 8,
  "dataStorageSizeInTBs" : 1,
  "source" : "DATABASE",
  "sourceId" : "ocid1.autonomousdatabase.oc1.phx..<unique_ID>",
  "cloneType" : "FULL"
}
```

Reference:

<https://docs.cloud.oracle.com/iaas/api/#/en/database/20160918/AutonomousDatabase/CreateAutonomousDatabase>

M09Q02 Which is the correct subset of services offered via OCI-CLI (command line interface) for Autonomous Database (ADB) via calls made to the OCI API's?

- A. Start, Delete, Update, Query, Stop
- B. Create, Query, Update, List, Start
- C. Create, Query, List, Stop, Restore
- D. Create, Get, List, Stop, Restore

Answer: D

Explanation:

A, B & C are wrong as there is no CLI command named as query for Autonomous Database

Autonomous Database – Using OCI-CLI Supported Services

Autonomous Database		Autonomous Database
		Backup
– Create	Creates a new Autonomous Database.	– Create
– Delete	Deletes the specified Autonomous Database	– Get
– generate-wallet	Creates and downloads a wallet	– list
– Get	Gets the details of the specified Autonomous Database	
– List	Gets a list of Autonomous Database	
– Restore	Restores an Autonomous Database	
– Start	Starts the specified Autonomous Database	
– Stop	Stops the specified Autonomous Database	
– Update	Updates the specified Autonomous Database	

Reference:

https://docs.cloud.oracle.com/iaas/tools/oci-cli/latest/oci_cli_docs/cmdref/db.html

M09Q03 How can an Autonomous Database resource be provisioned without logging into the Oracle Cloud Infrastructure console?

- A. It cannot be done.
- B. Using the Oracle Cloud Infrastructure Command Line interface tool or REST API calls.
- C. Using Database Configuration Assistant (DBCA) on the database server.
- D. Connecting to the Cloud Infrastructure Command console via SSH wallet.

Answer: B

Explanation:

- A is wrong as it can be done using REST API's, OCI-CLI & Terraform
- B is correct as Using the Oracle Cloud Infrastructure Command Line interface tool or REST API calls
- C is wrong you cannot create Autonomous Database using DBCA
- D is wrong as we cannot connect to Cloud Infrastructure Command console via SSH wallet

Reference:

<https://docs.cloud.oracle.com/iaas/Content/Database/Tasks/adbcreating.htm>

M09Q04 Which open source orchestration tool can be used to provision autonomous database resources in Oracle Cloud Infrastructure?

- A. REST API
- B. Terraform
- C. Enterprise Manager
- D. Docker

Answer: B

Explanation:

- A is wrong as REST API is not open source, it is custom built
- B is correct as Terraform is an open source orchestration tool and can be used to provision autonomous database resources in OCI
- C is wrong as Enterprise Manager is not an open source orchestration tool and we cannot provision autonomous database resources in Oracle Cloud Infrastructure
- D is wrong as through docker we cannot provisioned Autonomous Database

Reference:

https://www.terraform.io/docs/providers/oci/r/database_autonomous_database.html

<https://blogs.oracle.com/cloud-infrastructure/getting-started-with-the-resource-manager-on-oracle-cloud-infrastructure>

Module 10: Migrating Oracle Database To Autonomous Database

M10Q01 What are two advantages of using Data Pump to migrate your Oracle Databases to Autonomous Database?

- A. Data Pump is platform independent - it can migrate Oracle Databases running on any platform.
- B. Data Pump is faster to migrate database than using RMAN.
- C. Data Pump creates the tablespaces used by your Autonomous Database.
- D. Data Pump can exclude migration of objects like indexes and materialized views that are not needed by Autonomous Database.

Answer: A, D

Explanation:

A is correct as Data Pump is platform independent - it can migrate Oracle Databases running on any platform

B is wrong though Data Pump is faster to migrate database than using RMAN but RMAN backups are not supported in Autonomous Database

C is wrong as Data Pump does not create any tablespace in Autonomous Database and also you are not allowed to create any Tablespace in Autonomous Database

D is correct the Data Pump uses exclude and data_options parameters ensure that the object types not required in Autonomous Database are not exported and table partitions are grouped together so that they can be imported faster during the import to Autonomous Database. Oracle recommends using the following Data Pump Export parameters for faster and easier migration:

– exclude=index, cluster, indextype, materialized_view, materialized_view_log,
materialized_zonemap, db_link data_options=group_partition_table_data parallel=n
schemas=schema name dumpfile=export%u.dmp

Reference:

<https://docs.oracle.com/en/cloud/paas/atp-cloud/atpug/load-data-data-pump.html#GUID-297FE3E6-A823-4F98-AD50-959ED96E6969>

M10Q02 A customer wants to migrate to Autonomous Database (ADB) but only allows for a very small window of downtime. Golden Gate was advised to be used during the migration. For maximum reassurance of their end-users, the customer also would like to use Golden Gate as a fall-back scenario for the first 6 months after the migration. If customers complain, the on-premise data can be synchronised with the ADB instance for a switch back. Which statement about the migration using Golden Gate is correct?

- A. Migration to ADB is not possible using Golden Gate because the apply-process cannot be installed on ADB
- B. The fallback scenario is not possible using Golden Gate because the capture-process cannot be installed on ADB.
- C. Golden Gate on premise is not certified with ADB because Golden Gate Cloud Service exists for this.
- D. The described scenario is correct, can be used for migration and fallback scenarios.
- E. Only the migration to ADB is possible from an on-premise installation of Golden Gate.

Answer: B

Explanation:

A is wrong as we can migrate to ADB using Golden Gate

C is wrong as Oracle GoldenGate On-Premises releases 12.3.0.1.2 and later as a source is supported

E is wrong as Sources for ADB GoldenGate target databases can be: – Oracle Database on-premises 12.3.0.1.2 – Oracle Database Cloud Service on Oracle Cloud – Oracle Exadata Cloud Service on Oracle Cloud and Oracle Cloud at Customer, also any supported release of Oracle GoldenGate for any supported database and operating system combination that can send trail data to Oracle GoldenGate for Oracle Database release 12.3.0.1.2 and later, can be used as source systems.

Reference:

<https://docs.oracle.com/en/middleware/goldengate/core/18.1/oracle-db/replicating-data-oracle-autonomous-data-warehouse-cloud.html#GUID-E878AFC9-40F4-4F5E-8A72-E1A50ECECC82>

M10Q03 Which method can be used to migrate on-premises databases to Autonomous Databases in cloud?

- A. RMAN backup & restore
- B. Original Import/Export tools
- C. Physical migration method like database cloning
- D. Data Pump

Answer: D

Explanation:

MIGRATION

How can I migrate my existing Oracle Database to Oracle Autonomous Database?

- 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, database clones, and transportable tablespaces to move your existing database to ADB.

Reference:

<https://k21academy.com/1z0-931/migrate-database-to-oracle-autonomous-database-cloud-options-data-pump-sql-loader-dbms-cloud-goldengate/>

<https://www.oracle.com/database/technologies/datawarehouse-bigdata/adb-faqs.html#MIGRATION>

M10Q04 Migrating an on-premise database to Autonomous Database (ADB) for large amounts of data involves multiple steps like creating a credential object, creating (access to) a storage object/location, running data-pump export and running a data-pump import. Which three statements are true for SQL Developer (18.3 and up) in combination with ADB Data Loading?

- A. SQL Developer can only export/move/import files using data-pump from databases running on Linux systems.
- B. SQL Developer can import .csv files into ADB which are located on the system where SQL Developer is running.
- C. SQL Developer can be started from the ADB Cloud console but only for data loading scenarios.

D. SQL Developer can be used to export/move/import of a database to ADB in I set of wizard steps.

E. SQL Developer can import files (.dmp and .csv for example) into ADB which are located on Amazon S3 Object storage.

Answer: B, D, E

Explanation:


A is wrong as Data pump is platform independent and you can import from database running on any OS.

B is correct as SQL Developer can import .csv files into ADB which are located on the system where SQL Developer is running

C is wrong as SQL Developer can be started from the ADB Cloud but can be used for other purposes as well such as data loading, running queries against autonomous database etc

D is correct as SQL Developer can be used to export/move/import of a database to ADB in I set of wizard steps.

E is correct as Developer can import files (.dmp and .csv for example) into ADB which are located on Amazon S3 Object storage



Reference: <https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/sql-developer-web.html#GUID-102845D9-6855-4944-8937-5C688939610F>

<https://www.thatjeffsmith.com/archive/2018/07/loading-data-from-oss-to-oracle-autonomous-cloud-services-with-sql-developer/>

M10Q05 The 3rd party application that your customer wants to migrate to Autonomous Database (ADB) has some specific demands like tablespace names, usernames and init.ora parameters. The decision was made to adhere to the suggested migration method using an instant client and the data-pump version that was suggested (and came with it). Which statement is true about the migration of the application's database success?

- A. The suggested data-pump version will create an alias for non-standard tablespace names so the migration is successful.
- B. The migration can be technically a success but the 3-rd party vendor needs to support the result.

- C. The tablespace names will result in a blocking error during data-pump import because of ADB limitations.
- D. The migration can be a success, both technically and functional due to data-pump enhancements.

Answer: D

Explanation:

Storage attributes for tables are ignored

Reference:

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/load-data-data-pump.html#GUID-9F6B551A-9DB7-44BE-A8ED-08E300C56DB1>

M10Q06 Which statement about the Export Wizard used to export database objects and data is NOT correct?

- A. Export DDL includes features such as "Show schema," "Storage," and "Terminator."
- B. If "Clipboard" is selected as the "Output," the output will be placed on the system clipboard, so that it can be pasted into a file, a command line, or other location appropriate for the format.
- C. If "Dependents" is checked as a DDL Option, for non-privileged users, only dependent objects in their schema are exported.
- D. If "Grants" is checked as a DDL Option, GRANT statements are Included for any grant objects on the exported objects, including those owned by the SYS schema.

Answer: D

Explanation:

A is correct as Export DDL includes features such as "Show schema," "Storage," and "Terminator."

B is correct as If "Clipboard" is selected as the "Output," the output will be placed on the system clipboard, so that it can be pasted into a file, a command line, or other location appropriate for the format.

C is correct as If "Dependents" is checked as a DDL Option, for non-privileged users, only dependent objects in their schema are exported.

D is incorrect as this is a security violation.

Reference:

https://docs.oracle.com/cd/E17781_01/server.112/e18804/impexp.htm

<https://www.thatjeffsmith.com/archive/2012/06/copying-columns-from-grid-to-clipboard-in-sql-developer/>

M10Q07 What are the two methods that could be used during the migration of your existing Oracle database to Autonomous Database?

- A. Recovery Manager (RMAN)
- B. Data Pump
- C. Golden Gate
- D. CSV files copied to Autonomous Database block storage

Answer: B, C

Explanation:

MIGRATION

How can I migrate my existing Oracle Database to Oracle Autonomous Database?

- 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, database clones, and transportable tablespaces to move your existing database to ADB.

Reference:

https://k21academy.com/1z0-931/migrate-database-to-oracle-autonomous-database-cloud-options-data-pump-sql-loader-dbms_cloud-goldengate/

<https://www.oracle.com/database/technologies/datawarehouse-bigdata/adbfacts.html#MIGRATION>

M10Q08 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. Unplug to an xml file so database files and xml file can be uploaded to object storage.
- C. Unplug into a pdb archive which can be uploaded to object storage.

support@k21academy.com

<https://k21academy.com/autonomous03>

D. Create a database link from the source database to the ADB environment and clone the PDB.

Answer: A

Explanation:

unplug, plug OR Clone is not supported in ADB

Module 11: Autonomous Database For Experienced/Advanced Users

M11Q01 Which statement is true in regards to database links?

- A. You can call PL/SQL procedures and functions using a database link.
- B. Connect to Autonomous Database from remote database using a database link.
- C. Create a database link from one Autonomous Database to another Autonomous Database instance.
- D. Connect from Autonomous Database to remote database using a database link.

Answer: B

Explanation:

A is wrong as Calling PL/SQL programs using database links is not supported
B is correct as Connect to Autonomous Database from remote database using a database link is supported
C & D is also correct as Creating database link from one Autonomous Database to another Autonomous Database instance. & Connect from Autonomous Database to remote database using a database link is also supported now, wherein earlier it was not supported, so in terms of exam perspective B will be the right answer.

Reference:

<https://blogs.oracle.com/datawarehousing/how-to-create-a-database-link-from-an-autonomous-data-warehouse-to-a-database-cloud-service-instance-v2>

<https://blogs.oracle.com/datawarehousing/making-database-links-from-adw-to-other-databases>

M11Q02 Which three statements are true about procedures in the DBMS_CLOUD package?

- A. 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 bad file table on Autonomous Data Warehouse.
- B. The DBMS_CLOUD.PUT_OBJECT procedure copies a file from Cloud Object Storage to the Autonomous Data Warehouse.
- C. 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.
- D. The DBMS_CLOUD.CREATE_CREDENTIAL procedure stores Cloud Object Storage credentials in the Autonomous Data Warehouse database.
- E. The DBMS_CLOUD.DELETE_FILE procedure removes the credentials file from the Autonomous Data Warehouse database.

Answer: A, C, D

Explanation:

A is correct as 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.

B is wrong as The DBMS_CLOUD.PUT_OBJECT 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).

C is correct as The DBMS_CLOUD.CREATE_EXTERNAL_TABLE procedure creates an external table on files in the Cloud. This allows you to run queries on external data from Autonomous Data Warehouse.

D is correct as The DBMS_CLOUD.CREATE_CREDENTIAL procedure stores Cloud Object Storage credentials in the Autonomous Data Warehouse database. ‘

E is wrong as DBMS_CLOUD.DELETE_FILE This procedure removes the specified file from the specified directory on Autonomous Data Warehouse.

Reference:

<https://docs.oracle.com/en/cloud/paas/autonomous-data-warehouse-cloud/user/dbms-cloud.html#GUID-3D0A5361-672C-467D-AA90-656B1B4D4E37>

Module 12: Autonomous Database Use Cases

M12Q01 Which two system privileges does a user need to create analytic views?

- A. CREATE ANALYTIC LEVEL
- B. CREATE ATTRIBUTE DIMENSION
- C. CREATE ANALYTIC HIERARCHY
- D. CREATE ANALYTIC VIEW
- E. CREATE ANALYTIC MEASURE

Answer: B, D

Explanation:

A is wrong as there is no system privilege named as CREATE ANALYTIC LEVEL

B is correct as CREATE ATTRIBUTE DIMENSION is a system privilege which allow you to Create an attribute dimension in the grantee's schema.

C is wrong as there is no system privilege named as CREATE ANALYTIC HIERARCHY, it is CREATE HIERARCHY

D is correct as CREATE ANALYTIC VIEW allow you to Create an analytic view in the grantee's schema.

E is wrong as there is no system privilege named as CREATE ANALYTIC MEASURE.

Example 24-1 Granting System Privileges

The following statements grant the CREATE system privilege to the user av_user.

```
GRANT CREATE ATTRIBUTE DIMENSION TO av_user;  
GRANT CREATE HIERARCHY TO av_user;  
GRANT CREATE ANALYTIC VIEW TO av_user;  
GRANT SELECT ANY TABLE TO av_user;
```

Example 24-2 Granting Object Privileges

Reference:


<https://docs.oracle.com/en/database/oracle/oracle-database/12.2/dwhsg/overview-analytic-views.html#GUID-6F948948-6AE6-4A89-8AAC-5B8952CEF41D>

M12Q02 Which operating system can Data Visualization Desktop be run on?

- A. Linux
- B. Solaris
- C. Windows
- D. AIX

Answer: C

Explanation:



Middleware / Technical Details / Oracle Data Visualization Desktop Download

Technical Details ▾ Downloads →

Download	Product	Version
Download	Oracle Data Visualization Desktop	12c (12.2.5.3.0)

Download Details

Download Instructions

The above download links will direct you to the Oracle eDelivery web site where you will complete your download. You will be prompted to login with your Oracle Single Sign On account, select a platform to download, and accept the appropriate license agreement to continue. Any additional technical resources, documentation, community resources, and related product information can be found below if applicable. For more information about eDelivery or if you find yourself unfamiliar with the new site, please take a moment to walk through the Oracle Software Delivery Cloud Upgrade Tour in the FAQ.

[eDelivery Frequently Asked Questions →](#)

Release Notes

Mac users:
If you're using a Mac and have Oracle Data Visualization Desktop installed, then you must uninstall it before installing the new version. See [Known Issues](#).

Academy link:
If you can't download samples from the Academy Link, then download samples from the Analytics Library.

System requirements:

- Operating System: [Microsoft Windows x64](#) (64-bit) 7 SP1+, 8.1, or 10; Windows Server 2012 R2; Sierra (10.12); High Sierra (10.13)
- CPU: Intel(R) Core(TM)2 Duo CPU E8400 @ 3.00GHz, 2992 Mhz 2 Cores, 2 Logical Processors or faster
- Memory: 4.00 GB Memory or more
- Minimum free disk space: 2GB plus space for any uploaded data files
- User privileges: You must have Admin privileges to perform the installation

For more information about installation, see [Data Visualization Desktop Known Issues](#) and the [Frequently Asked Questions](#) appendix of the Data Visualization Desktop User's Guide.

References:

<https://www.oracle.com/middleware/technologies/oracle-data-visualization-desktop.html>

M12Q03 Which two PL/SQL functions can be used to validate an analytic view?

- A. VALIDATE_HIERARCHY
- B. VALIDATE_LEVELS
- C. VALIDATE_ANALYTIC_VIEW
- D. VALIDATE_DIMENSION
- E. VALIDATE_MEASURES

Answer: A, C

Explanation:

PL/SQL Package for Analytic Views

You can validate the data for analytic view and hierarchy objects with the following procedures in the DBMS_HIERARCHY package:

- CREATE_VALIDATE_LOG_TABLE procedure
- VALIDATE_ANALYTIC_VIEW function
- VALIDATE_CHECK_SUCCESS function
- VALIDATE_HIERARCHY function

Reference:

<https://docs.oracle.com/en/database/oracle/oracle-database/19/dwhsg/overview-analytic-views.html#GUID-FC745EA1-721D-48E3-96D1-B4DE1758AFF8>

M12Q04 Which three data dictionary views contain information about analytic view objects?

- A. ALL_ANALYTIC_VIEW_ID_ATTRS
- B. ALL_ANALYTIC_VIEW_KEYS
- C. ALL_ANALYTIC_VIEW_LVLGRPS
- D. ALL_ANALYTIC_VIEW_PATHS
- E. ALL_ANALYTIC_VIEW_DIM_CLASS

Answer: B, C, E

Explanation:

Data Dictionary Views for Analytic Views

The following data dictionary views contain information about analytic view objects. Only the views with the prefix ALL are listed. Each view also has a corresponding DBA and USER version.

Analytic View Views

- ALL_ANALYTIC_VIEW_ATTR_CLASS
- ALL_ANALYTIC_VIEW_BASE_MEAS
- ALL_ANALYTIC_VIEW_CALC_MEAS
- ALL_ANALYTIC_VIEW_CLASS
- ALL_ANALYTIC_VIEW_COLUMNS
- **ALL_ANALYTIC_VIEW_DIM_CLASS**
- ALL_ANALYTIC_VIEW_DIMENSIONS
- ALL_ANALYTIC_VIEW_HIER_CLASS
- ALL_ANALYTIC_VIEW_HIERS
- ALL_ANALYTIC_VIEW_KEYS
- ALL_ANALYTIC_VIEW_LEVEL_CLASS
- ALL_ANALYTIC_VIEW_LEVELS
- ALL_ANALYTIC_VIEW_LVLGRPS
- ALL_ANALYTIC_VIEW_MEAS_CLASS
- ALL_ANALYTIC_VIEWS

Reference:

<https://docs.oracle.com/en/database/oracle/oracle-database/19/dwhsg/overview-analytic-views.html#GUID-FC745EA1-721D-48E3-96D1-B4DE1758AFF8>

M12Q05 When you connect Oracle Analytics Cloud to the Autonomous Data Warehouse, what file needs to be uploaded?

- A. OJDBC.PROPERTIES
- B. CWALLET.SSO
- C. TNSNAMES.ORA
- D. SQLNET.ORA

Answer: B

Explanation:

Connecting To Oracle Analytics Cloud



The screenshot shows the 'Create Connection' window in Oracle Analytics Cloud. It includes a title bar, a header with the Oracle logo, and a list of fields: 'Connection Name' (TestConnection), 'Description' (For Testing Purpose), 'Client Credentials' (cwallet.sso, highlighted with a yellow box), 'Username' (admin), 'Password' (masked with asterisks), and 'Service Name' (demo12_high). A 'Select...' button is next to the 'Client Credentials' field.

Reference:

<https://docs.oracle.com/en/cloud/paas/analytics-cloud/acubi/create-connections-oracle-adw.html>

M12Q06 Once you have a connection to the Autonomous Data Warehouse, how do you import tables into the Oracle Analytics Cloud (OAC)?

- A. Load data using Data Pump
- B. Create a Data Set
- C. Import Table
- D. Replicate the table using Golden Gate

Answer: B

Reference: <https://blogs.perficient.com/2019/07/22/loading-data-into-oracle-autonomous-data-warehouse-using-oac/>