

Oracle 1Z0-931-26

Oracle Autonomous AI Database Professional

For More Information – Visit link below:

<https://www.examsempire.com/>

Product Version

- 1. Up to Date products, reliable and verified.**
- 2. Questions and Answers in PDF Format.**



<https://examsempire.com/>

Visit us at: <https://www.examsempire.com/1z0-931-26>

Latest Version: 6.0

Question: 1

Which three operations are available for a quarterly maintenance update on Autonomous Container Database? (Choose three.)

- A. Patch immediately
- B. Reschedule and patch immediately
- C. Reschedule patching
- D. Skip patching
- E. Apply a partial patch

Answer: B, C, D

Explanation:

Oracle Autonomous Container Database (ACD) provides specific options for managing quarterly maintenance updates to ensure minimal disruption and flexibility. The three correct operations available are:

Reschedule and patch immediately (B): This option allows users to adjust the maintenance window to a more convenient time and apply the patch right away within that rescheduled window. It's useful when the default schedule conflicts with business operations but immediate patching is still desired.

Reschedule patching (C): This option enables users to defer the patching to a later maintenance window that suits their operational needs, providing flexibility without applying the patch immediately.

Skip patching (D): Users can choose to skip a quarterly maintenance update entirely, which is beneficial if the current database version is stable and no immediate updates are required. Oracle allows skipping up to two consecutive quarterly updates for Autonomous Container Databases on dedicated infrastructure.

The incorrect options are:

Patch immediately (A): This is not a standalone option for quarterly updates on ACDs. Patching occurs within scheduled maintenance windows, and immediate patching outside of rescheduling is not supported as a distinct choice.

Apply a partial patch (E): Oracle does not support applying partial patches during quarterly maintenance updates. Updates are delivered as complete bundles to ensure consistency and security.

This aligns with Oracle's maintenance policies for Autonomous Databases on dedicated infrastructure, where control over scheduling and skipping is provided, but partial patching is not an option.

Question: 2

Your customer wants to permanently disable scheduled maintenance on Autonomous Dedicated Infrastructure to keep their current database version. How can they achieve this?

- A. Change the Automatic Maintenance Schedule to 'No Preference'
- B. Change the Automatic Maintenance Schedule to 'None'
- C. You cannot permanently disable scheduled Automatic Maintenance
- D. Change the Automatic Maintenance Schedule to 'Specify a Schedule' but do not select any month, week, or day

Answer: C

Explanation:

Oracle Autonomous Database on Dedicated Infrastructure enforces automatic maintenance to ensure security, stability, and performance. The correct answer is:

You cannot permanently disable scheduled Automatic Maintenance (C): Oracle mandates periodic maintenance updates (e.g., patching) to keep the database secure and compliant with the latest fixes. Customers can influence the timing of these updates but cannot disable them permanently. This is a design principle of the Autonomous Database service to reduce administrative overhead while maintaining system integrity.

The incorrect options are:

Change the Automatic Maintenance Schedule to 'No Preference' (A): This setting allows Oracle to determine the maintenance window but does not disable maintenance. It still occurs as scheduled by Oracle.

Change the Automatic Maintenance Schedule to 'None' (B): There is no 'None' option in the maintenance scheduling settings for Autonomous Dedicated Infrastructure. This is not a valid configuration.

Change the Automatic Maintenance Schedule to 'Specify a Schedule' but do not select any month, week, or day (D): This is not a functional workaround. Specifying a schedule requires selecting a valid time window, and leaving it blank does not prevent maintenance; it simply reverts to Oracle's default scheduling.

Oracle's documentation confirms that while customers can reschedule or skip specific maintenance runs (up to two consecutive quarters), permanently disabling automatic maintenance is not permitted to ensure the system remains up-to-date and secure.

Question: 3

Which statement is false about Autonomous Database Oracle Client Credentials (Wallets)?

- A. The Wallet for the Autonomous Database is the same as the TDE wallet
- B. You must have an Oracle Client Credential wallet in order to connect to the Autonomous Database
- C. In addition to the Oracle Client Credential Wallet, a user must have a username and password in order to connect to the Autonomous Database
- D. The Oracle Client Credential file is downloaded as a ZIP file

Answer: A

Explanation:

Oracle Client Credentials (wallets) are essential for secure connections to the Autonomous Database. The false statement is:

The Wallet for the Autonomous Database is the same as the TDE wallet (A): This is incorrect. The Oracle Client Credential wallet (used for client authentication and connection security) is distinct from the Transparent Data Encryption (TDE) wallet (used to encrypt data at rest within the database). The client wallet contains certificates and keys for mutual TLS (mTLS) authentication, while the TDE wallet manages encryption keys for data security. They serve different purposes and are not the same.

The correct statements are:

You must have an Oracle Client Credential wallet in order to connect to the Autonomous Database (B): True. The wallet is required for mTLS, which is the default authentication method for Autonomous Databases, ensuring secure connections.

In addition to the Oracle Client Credential Wallet, a user must have a username and password in order to connect to the Autonomous Database (C): True. Alongside the wallet, a database username and password are needed for full authentication (mTLS + user credentials).

The Oracle Client Credential file is downloaded as a ZIP file (D): True. When downloaded from the Oracle Cloud Infrastructure (OCI) console, the wallet is provided as a ZIP file containing configuration files like tnsnames.ora and certificates.

This distinction is critical for understanding security mechanisms in Autonomous Database.

Question: 4

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

- A. Using the DBCA on the database server
- B. Connecting to the cloud infrastructure console using the SSH wallet
- C. It cannot be done
- D. Using the cloud infrastructure command line interface or REST API calls

Answer: D

Explanation:

Provisioning an Autonomous Database without using the OCI Console is possible through programmatic methods. The correct answer is:

Using the cloud infrastructure command line interface or REST API calls (D): The Oracle Cloud Infrastructure Command Line Interface (OCI CLI) and REST APIs allow users to provision and manage Autonomous Database resources programmatically. This method is ideal for automation or when GUI access is not preferred. For example, the OCI CLI command `oci db autonomous-database create` can be used to provision a database by specifying parameters like

compartment ID, database name, and workload type. Similarly, a REST API POST request to /autonomousDatabases achieves the same result.

The incorrect options are:

Using the DBCA on the database server (A): The Database Configuration Assistant (DBCA) is a tool for on-premises Oracle databases, not for cloud-based Autonomous Databases, which are fully managed by Oracle.

Connecting to the cloud infrastructure console using the SSH wallet (B): SSH wallets are for secure shell access to compute instances, not for provisioning databases or interacting with the OCI Console.

It cannot be done (C): This is false, as programmatic provisioning via CLI or API is explicitly supported.

This capability enhances automation and integration into DevOps workflows.

Question: 5

Which three functions are provided by Spatial Studio? (Choose three.)

- A. Map visualization
- B. Custom SQL queries
- C. Spatial analysis
- D. Geocoding
- E. Spatial data editing

Answer: A, C, E

Explanation:

Oracle Spatial Studio is a self-service tool for working with spatial data in Autonomous Database.

The three correct functions are:

Map visualization (A): Spatial Studio provides robust capabilities to visualize spatial data on interactive maps, enabling users to explore geographic patterns and relationships visually.

Spatial analysis (C): It offers tools for performing spatial operations like proximity analysis, spatial joins, and buffering, which are essential for deriving insights from geographic data.

Spatial data editing (E): Users can edit spatial data, such as modifying geometries or updating attributes, directly within Spatial Studio, making it a powerful tool for data management.

The incorrect options are:

Custom SQL queries (B): While Spatial Studio supports spatial operations, it is primarily a graphical tool and does not focus on executing custom SQL queries. Such functionality is more aligned with tools like SQL Developer.

Geocoding (D): Geocoding (converting addresses to coordinates) is not a core feature of Spatial Studio. It focuses on visualization, analysis, and editing rather than address-to-coordinate conversion, which is typically handled by separate Oracle services or tools.

These functions align with Spatial Studio's purpose of simplifying spatial data management and analysis.

Question: 6

Which index type is designed to index catalog information?

- A. CONTEXT
- B. CTXCAT
- C. CTXRULE
- D. CONSTRAINS

Answer: B

Explanation:

Oracle provides specialized index types for different data indexing needs. The correct answer is: CTXCAT (B): The CTXCAT index type is specifically designed for indexing catalog information, such as product catalogs or structured data with short text fields. It supports fast queries on structured data combined with text search, making it ideal for applications like e-commerce catalogs.

The incorrect options are:

CONTEXT (A): The CONTEXT index is used for full-text search on large unstructured text data (e.g., documents), not specifically for catalog information.

CTXRULE (C): The CTXRULE index is designed for rule-based classification of text, not for catalog indexing.

CONSTRAINS (D): This appears to be a typo (likely meant "CONSTRAINTS"), but constraints are not index types; they enforce data integrity rules, not indexing.

CTXCAT enhances performance for catalog-style queries, distinguishing it from other text index types.

Question: 7

Which option describes the connection types that third-party software can use to connect to an Autonomous Database (ADB)?

- A. New ADB-specific protocols
- B. Any protocol supported by Oracle Net Services
- C. SQL*NET only
- D. JDBC only

Answer: B

Explanation:

Third-party software can connect to an Autonomous Database using standard Oracle connectivity options. The correct answer is:

Any protocol supported by Oracle Net Services (B): Autonomous Database leverages Oracle Net Services (also known as SQL*Net) for connectivity, supporting protocols like TCP/IP with TLS for secure communication. This allows a wide range of third-party tools (e.g., SQL Developer, Tableau, or custom applications) to connect using established Oracle networking standards.

The incorrect options are:

New ADB-specific protocols (A): There are no proprietary, ADB-specific protocols; it uses Oracle's existing Net Services framework.

SQL*NET only (C): While SQLNet is part of Oracle Net Services, the broader term "Oracle Net Services" encompasses all supported protocols, not just SQLNet.

JDBC only (D): JDBC is a specific Java-based API for database connectivity, not a protocol, and it's just one of many ways to connect via Oracle Net Services.

This flexibility ensures compatibility with diverse client applications.

Question: 8

When integrating object storage with Autonomous Database, how effectively does Oracle use the files in the object storage?

- A. Prune columns in parquet files
- B. Scan partition tables
- C. Scan hybrid partition tables
- D. Prune columns in CSV files

Answer: B

Explanation:

When Autonomous Database integrates with object storage (e.g., Oracle Cloud Infrastructure Object Storage), it optimizes data access. The correct answer is:

Scan partition tables (B): Oracle leverages partition pruning when scanning tables stored in object storage. This means it only scans the relevant partitions based on query predicates, significantly reducing I/O and improving performance. This is particularly effective for partitioned tables stored in formats like Parquet or ORC, which support partitioning.

The incorrect options are:

Prune columns in parquet files (A): While column pruning is possible with columnar formats like Parquet, the primary optimization highlighted in Oracle's documentation for object storage integration is partition pruning, not column pruning alone.

Scan hybrid partition tables (C): "Hybrid partition tables" is not a standard term in this context. Oracle supports external partitioned tables, but the focus is on scanning partitioned tables, not a hybrid-specific optimization.

Prune columns in CSV files (D): CSV files are row-based and do not natively support column pruning, making this less effective compared to partition scanning.

Partition pruning is a key optimization for external tables in object storage, enhancing query efficiency.

Question: 9

Which predefined role that exists in Autonomous Database includes common privileges that are used by a Data Warehouse developer? (Choose the best answer.)

- A. ADBDEV
- B. DWROLE
- C. ADMIN
- D. ADWC

Answer: B

Explanation:

Autonomous Database provides predefined roles tailored to specific use cases. The correct answer is:

DWROLE (B): The DWROLE predefined role is designed for Data Warehouse developers. It includes privileges commonly needed for data warehousing tasks, such as creating tables, views, and materialized views, as well as executing analytical queries. This role is optimized for Autonomous Data Warehouse (ADW) workloads.

The incorrect options are:

ADBDEV (A): There is no predefined ADBDEV role in Autonomous Database; this appears to be a fictional or misinterpreted role.

ADMIN (C): The ADMIN role is a superuser role with full database privileges, far exceeding the needs of a typical Data Warehouse developer and not tailored to that specific use case.

ADWC (D): This is not a predefined role; it might be a typo or confusion with ADW (Autonomous Data Warehouse), but no such role exists.

DWROLE is the best fit for a Data Warehouse developer's needs.

Question: 10

To whom, and in which order, are dedicated Exadata Infrastructure resources provisioned?

- A. Fleet Administrator, Autonomous Exadata Infrastructure -> Autonomous Container DB -> Database Administrator -> Autonomous DB
- B. Fleet Administrator, Autonomous Exadata Infrastructure -> Database Administrator -> Autonomous Container DB -> Autonomous DB
- C. Database Administrator, Autonomous Container DB -> Autonomous DB -> Fleet Administrator -> Autonomous Exadata Infrastructure
- D. Database Administrator, Autonomous Exadata Infrastructure -> Fleet Administrator -> Autonomous Container DB -> Autonomous DB

Answer: A

Explanation:

Provisioning dedicated Exadata Infrastructure resources for Autonomous Database follows a specific hierarchical order:

Fleet Administrator, Autonomous Exadata Infrastructure -> Autonomous Container DB -> Database Administrator -> Autonomous DB (A):

Fleet Administrator provisions Autonomous Exadata Infrastructure: The Fleet Administrator, responsible for infrastructure management, first sets up the Autonomous Exadata Infrastructure (AEI), which includes the physical Exadata hardware and networking configuration.

Fleet Administrator provisions Autonomous Container DB (ACD): Within the AEI, the Fleet Administrator creates the ACD, a container that hosts multiple Autonomous Databases.

Database Administrator provisions Autonomous DB: Finally, the Database Administrator (DBA) provisions individual Autonomous Databases within the ACD, configuring them for specific workloads (e.g., ATP or ADW).

The incorrect options are:

B: The Database Administrator cannot provision the ACD before the AEI is set up by the Fleet Administrator.

C and D: The Database Administrator does not provision the AEI or ACD; these are infrastructure-level tasks handled by the Fleet Administrator.

This order ensures proper separation of infrastructure and database management responsibilities.

Thank You for Trying Our Product

Special 16 USD Discount Coupon: NSZUBG3X

Email: support@examsempire.com

**Check our Customer Testimonials and ratings
available on every product page.**

Visit our website.

<https://examsempire.com/>