Download Oracle 1Z0-931-25 Exam Dumps Free

Shared by Stevenson on 17-06-2026

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

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

Options:

A- ALTER SESSION SET APPLICATION_CONTINUITY=TRUE
B- EXEC DBMS_CLOUD_ADMIN.ENABLE_APP_CONT('<service_name>');
C- ALTER DATABASE SET APPLICATION_CONTINUITY=TRUE
D- ALTER SYSTEM SET APPLICATION_CONTINUITY=TRUE

Answer:

B

Explanation:

Application Continuity ensures seamless failover for applications:

Correct Answer (B): EXEC DBMS_CLOUD_ADMIN.ENABLE_APP_CONT('<service_name>'); is the correct command for Autonomous Database on Shared Infrastructure. It enables Application Continuity for a specific service (e.g., tpurgent), leveraging the DBMS_CLOUD_ADMIN package designed for cloud-specific administration.

Incorrect Options:

A: ALTER SESSION only affects the current session, not persistent database-wide settings.

C, D: ALTER DATABASE and ALTER SYSTEM are not supported in Autonomous Database's managed environment; Oracle restricts such commands.

This ensures application resilience in a managed setup.

Question 2

Question Type: MultipleChoice

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 any more 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?

Options:

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

Answer:

A

Explanation:

The scenario requires a low-maintenance, high-performance database for an OLTP web application with occasional reporting. The correct answer is:

ATP using 'tpurgent' and 'high' TNS services to separate connection types (A): Autonomous Transaction Processing (ATP) is ideal here. It's a fully managed database optimized for OLTP workloads (e.g., scheduling service requests) with zero maintenance overhead---Oracle handles patching, backups, and tuning. ATP supports multiple connection services:

'tpurgent': Prioritizes low-latency, time-critical transactions (e.g., customer scheduling requests), ensuring fast response times for the web app.

'high': Suited for high-concurrency or reporting queries, allowing operational reports to run during non-peak hours without impacting the OLTP workload.

By using these predefined services, the architecture separates transactional and reporting workloads within the same database, meeting performance needs without additional administration. For example, the web app connects via tpurgent for real-time updates, while a reporting tool uses high for batch queries at night, leveraging ATP's auto-scaling if needed.

The incorrect options are:

Since the application needs to be highly available, it should be deployed on a Kubernetes Cluster (B): This misinterprets the question---it's about the database service, not the application deployment. Kubernetes is for container orchestration, not a database solution, and adds complexity counter to the low-maintenance goal.

It is best to build a separate data warehouse, and move the OLTP data on a nightly basis (C): While a separate Autonomous Data Warehouse (ADW) could handle reporting, it requires data movement (e.g., via ETL), increasing admin effort and complexity, which the DBA Manager wants to avoid. ATP can handle both workloads with proper service separation.

ADW since operational reporting is a higher priority in this scenario (D): ADW is optimized for analytics, not OLTP. The web app's transactional needs are primary, with reporting secondary and occasional, making ATP more suitable.

ATP's self-managing nature and service flexibility make it the best fit.

Question 3

Question Type: MultipleChoice

A Business Analyst joined your organization and wants to explore the database tools. When restoring or cloning an Autonomous Database (ADB), you must select a backup that is at least how old?

Options:

A- 24 hours
B- 5 minutes
C- 2 hours
D- 1 day

Answer:

A

Explanation:

Full Detailed In-Depth Explanation:

When restoring or cloning an Autonomous Database (ADB), Oracle enforces a minimum backup age to ensure data consistency and integrity. The official Oracle documentation specifies that backups used for these operations must be at least 24 hours old. This requirement exists because:

Backups need time to complete and stabilize, ensuring all transactions are fully committed and the backup is consistent.

Recent backups (e.g., less than 24 hours old) may still be in progress or lack full verification, risking incomplete or corrupted restores/clones.

Options B (5 minutes), C (2 hours), and D (1 day) are either too short or redundant:

5 minutes and 2 hours: Too recent, violating the 24-hour rule.

1 day: Matches 24 hours but is less precise than the explicit ''24 hours'' phrasing in the documentation.

For the Business Analyst's exploration, they can access tools like SQL Developer Web or Data Load via the OCI Console under the ''Tools'' tab, but this question focuses on the backup age constraint, making A the best answer.

Question 4

Question Type: MultipleChoice

Which statement is FALSE about the Catalog tool?

Options:

A- You can delete an entity from the Autonomous Database via the Catalog tool.
B- You can limit the display of entities to only those with certain names, schemas, types, or applications by entering values in the Search field.
C- You can view the data definition of an entity.
D- Impact shows the entities that depend on the source.

Answer:

A

Explanation:

The Catalog tool in Autonomous Database provides metadata exploration:

Correct Answer (A): ''You can delete an entity from the Autonomous Database via the Catalog tool'' is false. The Catalog is a read-only interface for viewing database objects (tables, views, etc.) and their relationships; it does not support modification or deletion actions.

True Statements:

B: Filtering by name, schema, type, or application is a core search feature.

C: Viewing DDL (data definition language) for entities is supported.

D: The Impact tab shows dependencies, aiding change analysis.

This ensures safe metadata exploration without risking data changes.

Question 5

Question Type: MultipleChoice

DBSRV1 is an Autonomous Database with the display name PRODDB. You plan to take manual backups of this database periodically. What must be the name of the bucket that will be used to store manual backups?

Options:

A- backup_proddb
B- You can use the Service Console to create manual backups of Autonomous Database.
C- backup_dbsrv1
D- Private endpoints can be configured when you clone a shared Autonomous Database that was initially created with public endpoints.

Answer:

D

Explanation:

Manual backups in Autonomous Database are stored in OCI Object Storage, and the bucket name follows a strict convention. The correct name is:

BACKUP_DBSRV1 (D): For manual backups, ADB uses a bucket named BACKUP_<DB_UNIQUE_NAME> (all uppercase). Here, DBSRV1 is the database's unique name (not the display name PRODDB), so the bucket must be BACKUP_DBSRV1. This bucket is automatically created by Oracle when you initiate a manual backup via the OCI console (e.g., ''Create Manual Backup'' button) or API (e.g., oci db autonomous-database backup create). For example, backing up DBSRV1 on March 18, 2025, stores files like backup_DBSRV1_20250318.dmp in this bucket. The unique name ensures uniqueness across the tenancy, as display names (e.g., PRODDB) might not be unique but the DB_UNIQUE_NAME is system-assigned and distinct.

The incorrect options are:

backup_proddb (A): This uses the display name proddb and lowercase, violating the BACKUP_<UNIQUE_NAME> uppercase convention.

BACKUP_PRODDB (B): While uppercase, it uses the display name PRODDB, not the unique name DBSRV1, so it's incorrect. Display names are for human readability, not system identification.

backup_dbsrv1 (C): This uses lowercase and lacks the exact BACKUP_ prefix format, failing Oracle's naming rule.

The rigid naming ensures consistency and automation in ADB's backup system.

Question 6

Question Type: MultipleChoice

You see a clock icon in the Status column on the SQL Monitoring tab of Performance Hub. What does it indicate?

Options:

A- The SQL statement is queued.
B- The SQL statement is executing.
C- The SQL statement did not complete either due to an error.
D- The SQL statement completed its execution.

Answer:

A

Explanation:

The Performance Hub in Autonomous Database provides real-time SQL monitoring with status indicators:

Correct Answer (A): A clock icon in the Status column indicates ''The SQL statement is queued.'' This means the statement is waiting in a queue (e.g., due to resource limits or consumer group settings) and has not yet started executing.

Incorrect Options:

B: An executing statement typically shows a green progress bar or running icon, not a clock.

C: An error would display a red icon or error symbol, often with details in the UI.

D: A completed statement shows a checkmark or similar success indicator.

This visual cue helps identify resource contention or scheduling delays.

Question 7

Question Type: MultipleChoice

What is a best practice when planning and instituting access controls for your Autonomous Dedicated environment regarding subnets, compartments, and user groups?

Options:

A- Only 1 subnet and 1 compartment are allowed, multiple groups highly advised.
B- Only 1 of each is allowed to be allocated per environment.
C- Create a separate VCN that contains only public subnets.
D- Create at least 2 of each resource.

Answer:

D

Explanation:

When designing access controls for an Autonomous Dedicated environment, Oracle recommends redundancy and separation for security and manageability:

Correct Answer (D): ''Create at least 2 of each resource'' (subnets, compartments, user groups) is a best practice:

Subnets: Using at least two subnets (e.g., one public, one private) enhances security by segregating traffic and provides failover options if one subnet encounters issues.

Compartments: Multiple compartments allow logical separation (e.g., dev, test, prod), simplifying access control and resource management.

User Groups: At least two groups (e.g., admins, developers) streamline permission assignments and reduce the risk of over-privileging users.

Incorrect Options:

A: OCI allows multiple subnets and compartments; restricting to one limits flexibility and security.

B: There's no such limitation; multiple resources are supported and encouraged.

C: A VCN with only public subnets contradicts security best practices, as private subnets are critical for database isolation.

This approach aligns with Oracle's guidance for secure and scalable deployments.

Question 8

Question Type: MultipleChoice

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

Options:

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

Answer:

D

Explanation:

SSH keys in OCI are used for secure access, but their role in API calls needs clarification. The correct answer is:

Navigate to Identity, select Users panel on the console and select 'Add Public Key' (D): This is the correct process, but with a caveat---it's about API authentication, not SSH for API calls per se. In OCI, API calls are authenticated using API keys (RSA key pairs), not SSH keys directly. To enable API access for a user, you generate a public/private key pair (e.g., using openssl), then add the public key in the OCI console:

Go to ''Identity & Security'' > ''Users.''

Select the user (e.g., john.doe).

Under ''API Keys,'' click ''Add API Key.''

Upload the public key (e.g., ~/.oci/oci_api_key_public.pem).

This associates the key with the user, allowing API calls (e.g., oci db autonomous-database list) authenticated via the private key and config file (e.g., ~/.oci/config). The question's phrasing (''SSH key'') likely misuses terminology, intending ''API key,'' as SSH keys are for compute instance access, not APIs. For example, a user might add a key to call the ADB API, securing requests with a signature.

The incorrect options are:

On the Autonomous Database Console (A): The ADB console manages database-specific settings (e.g., wallets), not user API keys, which are handled at the tenancy level under Identity.

SSH keys are not required in Oracle Cloud Infrastructure (B): False in general---SSH keys are needed for compute instances---but misleading here, as API calls use API keys, not SSH keys. Authentication (e.g., via tokens or keys) is required for APIs.

SSH keys cannot be added from console. They have to be added using REST APIs only (C): False. The console supports adding API keys under the Users panel; REST APIs (e.g., CreateApiKey) are an alternative, not the only method.

The correct path reflects OCI's user management for API access, despite the SSH terminology confusion.

Question 9

Question Type: MultipleChoice

As a database architect, you are tasked with configuring a high concurrency, production OLTP (EXAM) 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?

Options:

A- TPURGENT
B- TP
C- HIGH
D- MEDIUM

Answer:

C

Explanation:

Autonomous Transaction Processing (ATP) provides predefined connection services to optimize different workload types. For a high-concurrency OLTP application with parallel reporting queries, the appropriate service is:

HIGH (C): The HIGH connection service in ATP is designed for workloads requiring high concurrency and throughput, making it ideal for a production OLTP application like 'EXAM' that handles many simultaneous transactions (e.g., user requests). Additionally, HIGH supports parallel query execution, which is critical for running reporting queries efficiently during off-peak times. It allocates more resources (e.g., OCPUs and I/O) compared to other services, ensuring performance for both transactional and analytical tasks. For example, the OLTP app might use HIGH to process thousands of concurrent user requests, while a nightly report leverages parallel processing to aggregate data quickly. In the wallet's tnsnames.ora, the HIGH service might look like exam_high = (DESCRIPTION=...), offering the best balance for this mixed workload.

The incorrect options are:

TPURGENT (A): The TPURGENT service prioritizes low-latency, time-critical transactions (e.g., sub-second response times for single-user actions). It's optimized for latency-sensitive operations, not high concurrency or parallel reporting, and would underperform for the broader OLTP and reporting needs here. For instance, it's better suited for a single critical transaction than a multi-user system.

TP (B): The TP (Transaction Processing) service is a general-purpose option for OLTP workloads, offering moderate concurrency and latency. However, it's not optimized for high concurrency or parallel query execution, making it less suitable for a production app with reporting demands. It's a middle ground, not a top-tier performer like HIGH.

MEDIUM (D): The MEDIUM service balances performance and resource usage for moderate workloads. It supports some concurrency but lacks the resource allocation and parallel execution capabilities of HIGH, making it inadequate for a high-concurrency OLTP system with reporting requirements.

The HIGH service's ability to handle both high transactional volume and parallel reporting aligns perfectly with the scenario's strict performance needs, all within ATP's zero-maintenance framework.

Question 10

Question Type: MultipleChoice

Which statement is correct with respect to the required action to move Oracle Autonomous Database resources to a different compartment?

Options:

A- Moving an Autonomous Database instance does not include its automatic backups.
B- You do not require sufficient access permissions on the compartment that the resource is being moved to.
C- Autonomous Exadata Infrastructure instances and Autonomous Container Databases have no dependent resources that move with them. Associated (non-dependent) resources remain in their current compartments.
D- Moving the compartment of the Autonomous Database also moves the Autonomous Container Database and Autonomous Exadata Infrastructure.

Answer:

C

Explanation:

Moving resources between compartments in OCI has specific rules:

Correct Answer (C): ''Autonomous Exadata Infrastructure instances and Autonomous Container Databases have no dependent resources that move with them. Associated (non-dependent) resources remain in their current compartments'' is correct. For dedicated deployments, moving these resources leaves related items (e.g., backups, network configurations) in their original compartments unless explicitly moved.

Incorrect Options:

A: Backups move with the database instance in shared infrastructure, not dedicated.

B: Permissions are required for both source and target compartments.

D: Moving an Autonomous Database doesn't automatically move its container or infrastructure in dedicated setups.

This reflects OCI's compartmental resource management.