

Oracle Cloud Database Services 2026 Professional

Oracle 1z0-1093-25

Version Demo

Total Demo Questions: 10

Total Premium Questions: 106

Buy Premium PDF

<https://passqueen.com>

support@passqueen.com

QUESTION NO: 1

Which two statements are accurate regarding the placement of VM Clusters within an Exadata Cloud Infrastructure?

- A. Only one VM Cluster can exist within an Exadata Cloud Infrastructure.
- B. Multiple VM Clusters can be created within a single Exadata Cloud Infrastructure, allowing for workload isolation.
- C. VM Clusters can span multiple Exadata Cloud Infrastructures, providing a single, unified database environment.
- D. Each VM Cluster must be associated with its own dedicated VCN, isolating network traffic between clusters.
- E. VM Clusters within the same Exadata Infrastructure can share a single VCN, simplifying network management.

ANSWER: B E

Explanation:

An Exadata Cloud Infrastructure is designed to host more than one VM Cluster. This enables organizations to partition the physical Exadata resources into separately managed database environments, supporting workload separation, administrative isolation, and independent scaling or lifecycle operations for the VM Clusters. Each VM Cluster is deployed on a particular Exadata Cloud Infrastructure; it is not a distributed cluster extending across separate Exadata Cloud Infrastructures.

VM Clusters created on the same Exadata Cloud Infrastructure can use the same virtual cloud network (VCN). This allows customer networking to be organized under a shared VCN while VM Clusters retain their own network configuration details, including their relevant subnets and addresses. Using a common VCN can simplify connectivity design for application tiers and administrative access, provided subnet and security configuration are planned appropriately for the individual clusters.

Oracle documents that an Exadata Cloud Infrastructure can contain one or more VM Clusters and describes the VCN and subnet networking selected when provisioning a VM Cluster. See the [Oracle Exadata Cloud Infrastructure overview](#) and [Create a VM Cluster documentation](#).

QUESTION NO: 2

What is the primary purpose of installing and configuring the Oracle Management Agent when enabling Database Management Service for a cloud database on a Compute instance?

- A. To provide a secure communication channel between the Database Management Service and the database.
- B. To handle database backups and recovery operations.
- C. To manage the underlying operating system of the Compute instance.
- D. To automate database patching and upgrades.

ANSWER: A

Explanation:

To provide a secure communication channel between the Database Management Service and the database is correct because the Oracle Management Agent is the on-host component that connects a database running on an OCI Compute instance to OCI Database Management. After the appropriate Database Management plug-in is deployed and configured, the agent collects database metrics, configuration details, and performance information, then securely transmits that information to OCI services. This connection enables Database Management capabilities such as monitoring, performance analysis, SQL-related insights, and administration features for databases that are not managed directly by the OCI control plane.

The agent communicates outbound to Oracle Cloud Infrastructure service endpoints using authenticated, encrypted HTTPS connections, so it is not necessary to expose the database directly to the public internet for Database Management to obtain

its operational data. In practical terms, the agent serves as the secure intermediary between the database host and the Database Management service, allowing OCI to receive telemetry and issue supported management requests while maintaining controlled network connectivity. Oracle's setup guidance describes installing the Management Agent and deploying the Database Management plug-in as required steps for managing external databases, including databases hosted on Compute instances. See [Oracle Database Management documentation](#) and [Oracle Management Agent documentation](#).

QUESTION NO: 3

Which two metrics are available for monitoring within the Performance Hub's Real-Time SQL Monitoring feature in Database Management?

- A. Physical Reads
- B. Average Active Sessions
- C. Number of CPUs on the host
- D. Database Version

ANSWER: A B

Explanation:

Physical Reads and **Average Active Sessions** are available performance measurements in Database Management's Performance Hub context. Physical reads quantify database-block reads that require physical I/O, making the measure useful for identifying SQL executions with substantial storage activity. Real-Time SQL Monitoring exposes execution-level resource and progress information for qualifying long-running or parallel SQL, allowing administrators to investigate the I/O profile of an executing statement.

Average Active Sessions represents the workload load as sessions actively using CPU or waiting for non-idle database resources. Performance Hub uses this measure to present database activity over time and to correlate active workload with SQL activity, wait events, and resource consumption. This lets an administrator connect an observed rise in active demand to the SQL statements contributing to it, then drill into monitored execution details such as physical I/O. Oracle documents the SQL-monitoring data exposed through `V$SQL_MONITOR`, including physical read-related statistics, and documents Performance Hub as the Database Management interface for analyzing database activity and SQL performance.

References: [Oracle Database Reference: V\\$SQL_MONITOR](#) and [Oracle Cloud Infrastructure Database Management: Performance Hub](#).

QUESTION NO: 4

An Oracle NoSQL Database Cloud Service table is provisioned with insufficient capacity during a peak workload. Application logs show throttling for both read and write requests. Which two actions directly address the capacity limits involved? (Select two.)

- A. Increase the table's configured read throughput.
- B. Increase the table's configured write throughput.
- C. Increase the table's storage capacity without changing throughput settings.
- D. Add indexes without changing the table's throughput settings.
- E. Replace the application SDK without changing the table configuration.

ANSWER: A B

Explanation:

Oracle NoSQL Database Cloud Service controls provisioned table capacity separately for reads and writes. Increasing the table's read throughput addresses throttling caused by demand exceeding the configured read units. Increasing write throughput independently addresses throttling for write operations.

Increasing storage capacity does not increase the read or write request rate that the table can sustain. Adding indexes can improve some query access paths, but it does not directly raise the table's configured read or write throughput limits. Changing an SDK can affect application implementation but does not alter the service capacity allocated to the table.

QUESTION NO: 5

An analytics team wants to analyze historical files in OCI Object Storage together with operational data stored in MySQL tables, while avoiding a separate ETL process that loads every historical file into InnoDB. Which two statements accurately describe how HeatWave Lakehouse supports this requirement? (Select two.)

- A. HeatWave can query supported data files directly in object storage without first loading all of the data into MySQL tables.
- B. HeatWave can analyze eligible object-storage data together with data stored in MySQL tables.
- C. HeatWave automatically converts all object-storage files into InnoDB tables before executing a query.
- D. HeatWave Lakehouse continuously replicates MySQL data to a remote disaster-recovery region.
- E. HeatWave Lakehouse requires a separate data warehouse database for every object-storage data source.

ANSWER: A B

Explanation:

HeatWave Lakehouse enables HeatWave to query supported files in object storage directly. This allows data stored externally, such as CSV, Parquet, or Avro files, to be included in analytical processing without first copying all of that data into MySQL database tables.

Lakehouse queries can combine eligible object-storage data with data held in MySQL tables through the MySQL SQL interface. This supports analysis that joins operational and historical data while reducing unnecessary data movement and duplication. Lakehouse does not automatically convert object-storage files into InnoDB tables, and it is not a disaster-recovery replication mechanism for MySQL.

QUESTION NO: 6

When upgrading the Oracle Database software version on a BaseDB VM system, what is the primary responsibility of the database administrator (DBA) regarding the underlying operating system?

- A. The DBA is responsible for initiating and managing the operating system upgrade process through the Oracle Cloud Infrastructure console.
- B. The DBA must ensure the operating system is compatible with the target Oracle Database software version before initiating the database upgrade.
- C. The DBA has no direct responsibility for the operating system as it is fully managed by Oracle Cloud Infrastructure.
- D. The DBA should manually update the operating system packages after the database upgrade is successfully completed.

ANSWER: B

Explanation:

The DBA must ensure the operating system is compatible with the target Oracle Database software version before initiating the database upgrade. A database software upgrade is not solely a database-level activity: the target Oracle release must be certified and supported on the operating system release running on the Base Database Service VM. Verifying this prerequisite before starting protects the upgrade path from unsupported combinations and ensures that required operating-system capabilities, libraries, and platform dependencies are available for the target database home.

For BaseDB VM systems, Oracle provides managed database infrastructure, but the customer remains responsible for planning database lifecycle activities and validating prerequisites for a database upgrade. The DBA should review the applicable target-release requirements, confirm the current OS version is supported for that target database version, and perform any required OS lifecycle action through the supported service workflow before proceeding with the database software upgrade. This validation should be part of pre-upgrade planning alongside backups, compatibility checks, and upgrade readiness review. Oracle's Base Database Service documentation describes upgrade planning and service operations, while the database upgrade documentation details release and platform prerequisites. See [Oracle Base Database Service documentation](#) and [Oracle Database Upgrade Guide](#).

QUESTION NO: 7

A Base Database Service administrator wants to detect physical corruption in database files before a recovery event occurs. The administrator does not want to restore the database or create an additional backup merely to perform the check. Which RMAN operation is the most appropriate?

- A. Run RMAN VALIDATE DATABASE.
- B. Run RMAN RESTORE DATABASE.
- C. Run RMAN BACKUP DATABASE.
- D. Run RMAN RECOVER DATABASE.

ANSWER: A

Explanation:

RMAN VALIDATE DATABASE is the appropriate operation because it reads and validates database files for physical corruption without creating a backup. It can identify corrupt blocks early, allowing the DBA to investigate and plan remediation before corruption affects a recovery operation or production workload.

RESTORE DATABASE copies files from backup media and is used during recovery. BACKUP DATABASE creates a backup, while RECOVER DATABASE applies redo to restored datafiles. None of those operations is the focused integrity-validation operation required by the scenario.

QUESTION NO: 8

When enabling Database Management Service for a cloud database, which of the following is typically NOT a configurable parameter specific to the Database Management Service itself?

- A. The frequency at which performance metrics are collected.
- B. The specific metrics that are collected from the database.
- C. The maximum amount of storage space allocated for collected metrics.
- D. The database version.

ANSWER: D

Explanation:

The database version. is the correct choice because the database software version is an attribute of the cloud database deployment itself, selected during database system creation or changed later through a database upgrade process. Enabling Database Management associates the managed database with Oracle's management capabilities; it does not install, select, or modify the database release. The service discovers and uses the database's existing configuration and version information to provide functions such as monitoring, performance analysis, configuration insight, and administration features appropriate to the target database. Database version lifecycle operations therefore remain part of the database service's provisioning, maintenance, and upgrade workflows, rather than a setting of Database Management. Oracle documents database software selection as part of creating a database system, while its Database Management documentation

describes enabling management for an already deployed database. See Oracle's [database system creation documentation](#) and [Database Management enablement documentation](#).

QUESTION NO: 9

An Autonomous Database workload normally requires 4 OCPUs but experiences brief, unpredictable spikes that require approximately 10 OCPUs. The team wants to preserve the 4-OCPU baseline while avoiding permanent provisioning for peak capacity. Which configuration is the most appropriate?

- A. Increase the base OCPU count permanently to 10 OCPUs.
- B. Enable auto scaling while retaining a 4-OCPU base configuration.
- C. Enable Database Management to allocate additional OCPUs during periods of high load.
- D. Create a second Autonomous Database and use OCI Monitoring to distribute sessions between them.

ANSWER: B

Explanation:

Enabling auto scaling on the Autonomous Database is the correct choice. Auto scaling allows Autonomous Database to use additional CPU resources above its base OCPU count when workload demand requires it, then return to the configured baseline when the demand subsides. With a 4-OCPU base configuration, auto scaling can supply sufficient additional capacity for a short-lived 10-OCPU demand without permanently sizing the database for the peak.

Increasing the base OCPU count to 10 provides capacity but also permanently provisions resources that are not needed during normal operation. Database Management and OCI Monitoring can identify the spike, but neither service adds database CPU capacity. Creating a separate database would not transparently relieve the workload on the existing database.

QUESTION NO: 10

A BaseDB database requires a recovery strategy that provides rapid recovery from a complete DB system failure and retains a synchronized copy of the production database. Which two statements accurately describe the role of Oracle Data Guard in this design? (Select two.)

- A. A physical standby can receive and apply redo from the primary database to maintain a synchronized copy.
- B. A standby can be failed over to reduce recovery time after an unrecoverable primary DB system failure.
- C. Data Guard eliminates the need to maintain RMAN backups.
- D. Block change tracking performs the same role transition as a Data Guard failover.
- E. A Data Guard standby prevents logical application errors from being propagated to the standby.

ANSWER: A B

Explanation:

A physical standby database receives redo from the primary database and applies that redo to maintain a synchronized copy. Because the standby is already provisioned and maintained, it can assume the primary role through a failover when the original primary DB system is unavailable, producing a much lower recovery time than a restore-based recovery.

Data Guard complements rather than replaces RMAN backups. Backups remain necessary for recovery from problems such as logical corruption, user error, or other conditions that can be propagated to a standby. RMAN incremental backups and block change tracking improve backup efficiency, but they do not provide the role-transition capability of a synchronized standby database.