

# **Free Questions for 1Z0-076**

**Shared by Kelly on 04-10-2024**

**For More Free Questions and Preparation Resources**

**Check the Links on Last Page**

# Question 1

---

**Question Type:** MultipleChoice

---

Which two statements are true regarding asynchronous redo transport in a Data Guard

## Options:

---

- A-** This transport mode satisfies the minimum requirements for Maximum Availability data protection mode.
- B-** A transaction can commit without waiting for redo to be sent to any standby database in the data guard configuration.
- C-** This transport mode satisfies the minimum requirements for Maximum Performance data protection mode.
- D-** Real-time query performance on a physical standby database improves for current read requests when using this transport mode.
- E-** The performance of SQL apply on a logical standby database always improves when using this transport mode.

## Answer:

---

B, C

## Explanation:

---

Asynchronous redo transport is a method where the primary database does not wait for an acknowledgment from the standby database before committing transactions, which helps in minimizing the impact on the primary database's performance (B). This transport mode is associated with the Maximum Performance data protection mode, which prioritizes performance over synchronicity of data between the primary and standby databases (C). While it provides a level of data protection, there could be some data loss in the event of a primary database failure because redo data may not have been transmitted to the standby database at the time of the failure.

Reference: Oracle Data Guard Concepts and Administration documentation provides detailed explanations of different redo transport modes and their implications on data protection and performance. Asynchronous transport mode's behavior and association with Maximum Performance mode are outlined explicitly.

## Question 2

---

**Question Type:** MultipleChoice

---

Which three Data Guard monitoring activities may be performed using Enterprise Manager Cloud Control?

### Options:

---

- A- You can monitor the redo apply rate on a logical standby database.
- B- You can set a critical threshold on the redo generation rate metric for a primary database.

- C-** You can set a warning threshold on the redo generation rate metric for a physical standby database.
- D-** You can check if redo apply needs to be tuned.
- E-** You can check the potential data loss in the event of a disaster.
- F-** You can monitor the redo apply rate on a snapshot standby database.

### **Answer:**

---

A, B, E

### **Explanation:**

---

Enterprise Manager Cloud Control offers comprehensive monitoring capabilities for Oracle Data Guard environments. It enables monitoring the rate at which redo is being applied on a logical standby database (A), which is crucial for ensuring that the standby database is keeping up with the changes from the primary. It also allows setting thresholds on performance metrics, such as the redo generation rate on the primary database (B), to alert administrators when values exceed critical or warning thresholds. Additionally, it provides the capability to estimate the potential data loss in the event of a disaster (E), helping in disaster recovery planning and ensuring business continuity.

Reference: Oracle Enterprise Manager Cloud Control documentation provides extensive information on its monitoring features for Oracle Data Guard, including setting thresholds, estimating potential data loss, and tracking redo apply rates.

## Question 3

---

**Question Type:** MultipleChoice

---

Which TWO observations are true about the Far Sync instance?

### Options:

---

- A- Receives redo synchronously from the primary database
- B- Can be created using the RMAN DUPLICATE command
- C- Includes a standby control file, password file, data files, standby redo logs, and archive logs
- D- Can only be created using a series of SQL commands
- E- Applies redo received

### Answer:

---

A, E

### Explanation:

---

A Far Sync instance is a special kind of Oracle Data Guard configuration that allows synchronous redo transport from a primary database to a remote standby database with minimum impact on the primary database's performance. The Far Sync instance receives

redo data synchronously from the primary database (A), then ships it asynchronously to the remote standby database, thus extending zero data loss protection over longer distances and higher network latency environments than would be practical with a synchronous standby alone. The Far Sync instance does not apply the redo data; it just receives and ships it (E). A Far Sync instance does not have data files, and it cannot apply redo to stay synchronized with the primary database.

Reference: Oracle Database High Availability Overview and Oracle Data Guard Concepts and Administration documentation detail the role and configuration of Far Sync instances, including how they contribute to achieving zero data loss disaster recovery over long distances.

## Question 4

---

**Question Type:** MultipleChoice

---

On your logical standby database, you specified these rules:

```
SQL> EXECUTE DBMS_LOGSTDBY.SKIP (STMT => 'DML',-  
    SCHEMA_NAME => 'HR', -  
    OBJECT_NAME => 'EMP_NEW');
```

```
SQL> EXECUTE DBMS_LOGSTDBY.SKIP (STMT => 'DML',-  
    SCHEMA_NAME => 'HR', -  
    OBJECT_NAME => 'EMP_OLD');
```

After completion of the weekend batch cycle you attempt to delete the SQL Apply filters:

```
SQL> EXECUTE DBMS_LOGSTDBY.UNSKIP (STMT => 'DML',-  
    SCHEMA_NAME => 'HR', -  
    OBJECT_NAME => 'EMP%');
```

Which is TRUE regarding the execution of the UNSKIP procedure?

### Options:

---

- A-** It succeeds only if all DML statements executed on the primary have been applied on the logical standby deleting the SQL Apply filter.
- B-** It deletes both the SQL Apply filters.
- C-** It succeeds but the SQL Apply filters are not deleted.
- D-** It succeeds only if SQL apply is stopped before deleting the SQL Apply filter.
- E-** It returns an error because the syntax to delete a SQL Apply filter must specify the same object names as specified when the filter was added.

### Answer:

---

B

### Explanation:

---

The execution of the UNSKIP procedure is designed to remove SQL Apply filters that have been previously set up on a logical standby database. Based on the provided statements, the UNSKIP procedure is directed to delete any SQL Apply filters for DML statements associated with objects in the 'HR' schema that start with 'EMP'. Since both SKIP procedures had the same schema name ('HR') and statement type ('DML'), and the UNSKIP procedure uses a wildcard (%) for the object name, it will successfully remove both of the SQL Apply filters for 'EMP\_NEW' and 'EMP\_OLD', as both object names match the pattern provided in the UNSKIP procedure.

Reference: Oracle's Data Guard documentation and SQL Language Reference provide insights into managing SQL Apply filters on a logical standby database using the DBMS\_LOGSTDBY package. This includes adding and removing filters through SKIP and UNSKIP procedures.

## Question 5

---

**Question Type:** MultipleChoice

---

Which four statements are true regarding SQL Apply filters for a logical standby database?

### Options:

---

- A- They can be used to skip execution of DML triggers on a table while allowing the DML to execute.
- B- They can be used to skip CREATE TABLE commands.



- C-** They can be used to skip ALTER SYSTEM and ALTER DATABASE commands.
- D-** They can only be used to skip DML statements on a table.
- E-** They can be used to skip all SQL statements executed on a specific pluggable database (PDB) within a standby multitenant container database (CDB).
- F-** They can be used to stop SQL apply if it encounters an error.
- G-** They can be used to skip ALTER TABLE commands on specific tables.

**Answer:**

---

A, B, C, G

**Explanation:**

---

Based on the Oracle Database 19c documentation, the correct answers about SQL Apply filters for a logical standby database are: A. They can be used to skip execution of DML triggers on a table while allowing the DML to execute. B. They can be used to skip CREATE TABLE commands. C. They can be used to skip ALTER SYSTEM and ALTER DATABASE commands. G. They can be used to skip ALTER TABLE commands on specific tables.

Comprehensive Detailed Explanation: SQL Apply filters in a logical standby database can be set to control which SQL operations are applied to the standby. These filters allow for certain commands to be skipped, ensuring that they do not impact the standby database. For example, filters can be used to skip the execution of DML triggers to prevent them from firing during SQL Apply, while still allowing the underlying DML to be executed on the logical standby database. This is particularly useful when certain triggers are not desired to run in a standby environment. CREATE TABLE, ALTER SYSTEM, ALTER DATABASE, and specific ALTER TABLE commands can

also be skipped using SQL Apply filters to prevent unwanted structural changes or administrative operations from affecting the logical standby database. These capabilities provide a level of control to ensure that the logical standby database reflects only the desired state of the primary database.

Reference: Oracle Database SQL Language Reference and Oracle Data Guard Concepts and Administration guide offer comprehensive details on the use of SQL Apply filters, including the range of SQL statements that can be influenced by these filters in a logical standby database environment.

## Question 6

---

**Question Type: MultipleChoice**

---

Examine the Data Guard configuration:

```
DGMGRL> show configuration;
```

```
Configuration - Animals
```

```
Protection Mode: MaxAvailability
```

```
Databases:
```

```
dogs - Primary database
```

```
sheep - (*) Physical standby database
```

```
cats - Physical standby database
```

```
Fast-Start Failover: ENABLED
```

```
Configuration Status:
```

```
SUCCESS
```

What happens if you issue "switchover to sheep;" at the DGMGRL prompt?

### Options:

---

- A- The switchover succeeds and Cats becomes the new failover target.
- B- It results in an error indicating that a switchover is not allowed.
- C- The switchover succeeds but Dogs needs to be reinstated.
- D- The switchover succeeds and Fast-Start Failover is suspended.
- E- The switchover succeeds and Dogs becomes the new failover target.

## Answer:

---

E

## Explanation:

---

When issuing a 'switchover to sheep;' command in a Data Guard configuration, the primary database (Dogs) transitions to a standby role, and the target standby database (Sheep) becomes the new primary database. Fast-Start Failover (FSFO) remains enabled, but its target changes according to the new roles of the databases. Since Cats is also a physical standby database, it does not become the failover target by default unless it is specified in the broker configuration. After the switchover, the original primary (Dogs) becomes the new standby database and thus the new failover target for FSFO. Reference: Oracle Data Guard Broker documentation provides detailed procedures and explanations of switchover operations, including how FSFO targets are affected post-switchover. This behavior is consistent across different Oracle Database versions that support Data Guard and FSFO.

## Question 7

---

### Question Type: MultipleChoice

---

Examine the Data Guard configuration: DGMGRL> show configuration;

Configuration - Animals

Protection Mode: MaxPerformance

Databases:

dogs- Primary database

sheep - Physical standby database

cats- Snapshot standby database

Fast-Start Failover: DISABLED

Configuration Status: SUCCESS

You receive an error while attempting to raise the protection mode to Maximum Protection:

```
DGMGRL> edit configuration set protection mode as maxprotection;
```

Error: ORA-16627: operation disallowed since no standby databases would remain to support protection mode

Failed.

What can you conclude based on this error?

### **Options:**

---

**A-** The redo transport mode is set to async for the standby database Sheep.

**B-** The redo transport mode is set to asyn: for the standby database Cats.

**C-** The redo transport mode is set to async for both standby databases.

**D-** Cats is a snapshot standby database.

**D-** Cats is a snapshot standby database.

Comprehensive Detailed Explanation:

In an Oracle Data Guard environment, the Maximum Protection mode requires that all redo data be transmitted synchronously to at least one standby database, ensuring no data loss even in the event of a primary database failure. However, a snapshot standby database, by its nature, allows read-write access and is temporarily disconnected from the redo stream, which makes it unable to participate in the synchronous redo transport required by Maximum Protection mode. The presence of a snapshot standby database in the Data Guard configuration thus prevents the activation of Maximum Protection mode, as it cannot guarantee zero data loss without a standby database capable of receiving redo data synchronously.

Reference:

Oracle Data Guard documentation clearly outlines the requirements and restrictions of different protection modes, including the necessity for standby databases to participate in synchronous redo transport to enable Maximum Protection mode. The inability of snapshot standby databases to fulfill this requirement is a key consideration when planning Data Guard configurations and protection levels.

**Answer:**

---

D, D

**Explanation:**

---

The error indicates that switching the protection mode to Maximum Protection is not possible due to the presence of a snapshot standby database in the Data Guard configuration, which cannot participate in synchronous redo transport required by the Maximum Protection mode. Therefore, the correct answer is:

## Question 8

---

**Question Type:** MultipleChoice

---

Your Data Guard environment has a remote physical standby database with real-time query enabled, which is used for reporting, and a logical standby database used for DSS reporting.

Switchovers or failovers are possible due to testing or in case of a disaster.

Clients use local TNSNAMES.ORA files to define connection strings to the database instances.

Which three will prevent clients from connecting to the wrong database instances?

### Options:

---

**A-** Oracle Net connectivity to the primary database instance must be established on all the standby database instances.

- B-** The standby database services must be defined statically with the Listeners running on the standby database hosts.
- C-** The LOCAL\_LISTENER parameter on the primary database instance must always be set.
- D-** The client applications must use the correct TNS entries when requesting connections to the database instances.
- E-** Client TNS entries for the databases use the correct service names for the intended service.
- F-** The DB\_NAME and DB\_UNIQUE\_NAME parameters must be set to the same value for all the databases in the Data Guard environment.
- G-** A service name is registered with the local listener of each database instance.

**Answer:**

---

B, D, E

**Explanation:**

---

Based on Oracle Database 19c: Data Guard Administration documents, the three measures that can prevent clients from connecting to the wrong database instances during switchovers, failovers, or regular operations in a Data Guard environment are:

- B) The standby database services must be defined statically with the Listeners running on the standby database hosts.
- D) The client applications must use the correct TNS entries when requesting connections to the database instances.
- E) Client TNS entries for the databases use the correct service names for the intended service.



In an Oracle Data Guard configuration, correctly configuring Oracle Net Services (including TNS entries and listeners) is crucial for ensuring that clients connect to the appropriate database instance, whether it's the primary or standby. Defining services on the standby database and associating them with listeners ensures that client applications can connect to the standby when needed, especially useful in a role transition or when the standby is open for read-only access or real-time query. It's essential that TNS entries used by client applications specify the correct service names that correspond to the intended database roles, such as primary or standby. This setup facilitates seamless connectivity to the appropriate instance based on the role, especially critical during switchovers and failovers when the roles of the databases change.

Reference: Oracle's Data Guard concepts and administration guide provides extensive information on configuring network services for Data Guard environments, ensuring that applications connect to the correct database instance based on the current role of the databases in the Data Guard configuration.

## Question 9

---

**Question Type: MultipleChoice**

---

Your Data Guard environment contains a primary database and three standby databases with these attributes:

1. prod : Primary database
2. prod\_prq : Physical standby database with real-time query enabled used by reporting applications
3. prod\_lsby: Logical standby database used by DSS

#### 4. PROD\_SSBY: Snapshot standby database used for Real Application Testing

Which TWO can be used to prevent clients from connecting to the wrong database instance?

#### Options:

---

- A-** Create role based services with the `srvctl` utility when using clusterware for Oracle RAC databases or Oracle Restart for single instance Oracle databases.
- B-** Establish Oracle Net connectivity to the primary database instance from all the standby database instances.
- C-** Create a static service for each of the databases, register it with the local listener of each database instance, and add connection descriptors on clients to connect to those services.
- D-** Create database services on each of the standby databases, start the services, and add connection descriptors on the clients to connect to those services.
- E-** Create database services for each database and use event triggers to make sure that services are activated only when the database is in the correct role.

#### Answer:

---

D, E

#### Explanation:

---

Creating dedicated database services for each database instance (Option D) and utilizing event triggers to manage these services based on the role of the database (Option E) ensure that clients connect to the appropriate database instance based on its current role and state. This approach leverages the flexibility and control provided by Oracle Net services and database event management to direct client connections to the suitable primary or standby instance, enhancing the overall robustness and reliability of the Data Guard environment. Reference: Based on Oracle Database 19c best practices for managing connectivity and services in a Data Guard setup, including the use of role-based services and event-driven service management.

## Question 10

---

**Question Type: MultipleChoice**

---

Active Data Guard (ADG) databases are widely used to offload reporting or ad hoc query-only jobs from the primary database. Reporting workload profile is different from the primary database and often requires tuning.

Which tool is used to tune SQL workloads running on an ADG database?

### Options:

---

**A-** Standby Statspack

**B-** In-Memory Active Session History (ASH)

- C- Automatic Diagnostic Database Monitor (ADDM)
- D- Automatic Workload Repository (AWR)
- E- SQL Tuning Advisor

**Answer:**

---

D

**Explanation:**

---

AWR collects, processes, and maintains performance statistics for problem detection and self-tuning purposes. In an Active Data Guard environment, where the physical standby database can be used for read-only workloads, AWR can be instrumental in identifying performance bottlenecks and areas for optimization. It provides detailed reports that include wait events, time model statistics, and active session history, making it an invaluable tool for tuning SQL queries and overall database performance in an ADG setup.

## Question 11

---

**Question Type:** MultipleChoice

---

Your Data Guard environment consists of these components and settings:

1. A primary database
2. A remote physical standby database
3. Real-time query is enabled.
4. The redo transport mode is set to SYNC.
5. The protection mode is set to Maximum Availability.

You notice that queries executed on the physical standby database receive errors: ORA-03172: STANDBY\_MAX\_DATA\_DELAY of 15 seconds exceeded. Which two would you recommend to avoid this error?

### Options:

---

- A-** Increase the size of the buffer cache on the standby database instance.
- B-** Reduce I/O latency for the storage used by the primary database.
- C-** Increase the number of standby redo log files on the primary database.
- D-** Change the protection mode to Maximum Performance.
- E-** Increase the network bandwidth between the primary and standby databases.
- F-** Change the protection mode to Maximum Protection.

### Answer:

---

B, E

### **Explanation:**

---

The ORA-03172: STANDBY\_MAX\_DATA\_DELAY error indicates that the real-time query on the physical standby database is experiencing delays beyond the specified maximum data delay threshold. Increasing the network bandwidth (Option E) can enhance the speed at which redo data is transferred from the primary to the standby database, thereby reducing the likelihood of exceeding the STANDBY\_MAX\_DATA\_DELAY threshold. Reducing I/O latency on the primary database's storage (Option B) ensures that redo data is generated and shipped more efficiently, further mitigating the risk of delay. These actions, focused on optimizing data transfer and processing speed, address the root causes of the ORA-03172 error in a synchronous Data Guard configuration operating in Maximum Availability mode.

## **Question 12**

---

**Question Type:** MultipleChoice

---

Your Data Guard environment consists of these components and settings:

1. A primary database
2. Two remote physical standby databases
3. The redo transport mode is set to sync

4. Real-time query is enabled for both standby databases
5. The DB\_BLOCK\_CHECKING parameter is set to TRUE on both standby databases

You notice an increase in redo apply lag time on both standby databases.

Which two would you recommend to reduce the redo apply lag on the standby databases?

### Options:

---

- A-** Increase the size of standby redo log files on the standby databases.
- B-** Decrease the redo log file size on the primary database.
- C-** Increase the number of standby redo log files on the standby databases.
- D-** Lower DB\_BLOCK\_CHECKING to MEDIUM or low on the standby databases.
- E-** Increase the size of the buffer cache on the physical standby database instances.

### Answer:

---

A, D

### Explanation:

---

To reduce the redo apply lag on standby databases, one could increase the size of the standby redo log files. Larger redo log files can accommodate more redo data, which may reduce the frequency of log switches and allow for more continuous application of redo data.

Additionally, lowering the `DB_BLOCK_CHECKING` parameter to `MEDIUM` or `LOW` on the standby databases can help improve redo apply performance. High block checking can impose additional CPU overhead during the application of redo data, potentially increasing apply lag times. By reducing the level of block checking, you can lessen this overhead and help reduce the apply lag .



**To Get Premium Files for 1Z0-076 Visit**

**<https://www.p2pexams.com/products/1z0-076>**

**For More Free Questions Visit**

**<https://www.p2pexams.com/oracle/pdf/1z0-076>**

