Free Questions for 1Z0-076 by ebraindumps

Shared by Edwards on 24-05-2024

For More Free Questions and Preparation Resources

Check the Links on Last Page

Question 1

Question Type: MultipleChoice

You detected an unrecoverable archive gap in your data guard environment. So, you need to roll standby.

forward in time without applying a large number of archive log files using this command:

RMAN> RECOVER STANDBY DATABASE FROM SERVICE-;

When running this command, which of the following steps can be performed automatically?

1. Remember all data file names on the standby.

2. Restart standby in nomount.

3. Restore controlfile from primary.

4. Mount standby database.

5. Rename data files from stored standby names.

6. Restore new data files to new names.

7. Recover standby.

Options:

A- 2,3,5,6,7

B- 2,3,6,7

C- 1,3,5,6,7

D- 1,2,3,4,5,6,7

E- 1, 2,3,4,6,7

Answer:

E

Explanation:

The RECOVER STANDBY DATABASE FROM SERVICE command in RMAN is designed to automate various steps required to recover the standby database, especially when dealing with an archive gap. When this command is executed, the following actions can occur automatically:

Remember all data file names on the standby (1): RMAN has the capability to recall the names and paths of all data files associated with the standby database.

Restart standby in nomount (2): The standby database can be automatically restarted in the NOMOUNT state, allowing recovery operations to proceed without the database being open.

Restore controlfile from primary (3): RMAN can restore the control file from the primary database to the standby system, ensuring that the standby has the most up-to-date control file.

Mount standby database (4): After restoring the control file, the standby database is mounted to prepare for data file recovery.

Rename data files from stored standby names (5): Not typically done automatically by this command.

Restore new data files to new names (6): New data files added to the primary since the last synchronization can be restored to the standby with their correct names.

Recover standby (7): Finally, RMAN will apply any necessary redo logs to bring the standby database up to date with the primary.

While some steps, such as renaming data files (5), typically require manual intervention or scripting, most of the recovery process can be handled by RMAN automatically, streamlining the recovery of the standby database.

Oracle Database Backup and Recovery User's Guide

Oracle Data Guard Concepts and Administration Guide

Question 2

Question Type: MultipleChoice

Which three actions are performed by the START PLAN procedure of the DBMS ROLLING package?

Options:

A- converting the designated physical standby database into a logical standby database

B- creating a guaranteed restore point on the standby databases

C- building a LogMiner dictionary on the primary database instance

D- creating a guaranteed restore point on the primary database

E- starting media recovery on all the Leading Group Standby databases

F- switching the primary database to the logical standby role

Answer:

B, C, D

Explanation:

The DBMS_ROLLING package facilitates a rolling upgrade process across a Data Guard configuration. The START PLAN procedure in particular handles several critical actions, including:

Creating a guaranteed restore point on the standby databases (B): This ensures that the standby databases can be reverted to their state before the rolling upgrade process in case of any issues.

Building a LogMiner dictionary on the primary database instance (C): This is necessary for logical standby databases to interpret redo data during the SQL Apply process.

Creating a guaranteed restore point on the primary database (D): Similar to the standby databases, this ensures that the primary database can be reverted to a known good state if necessary. Reference:

Oracle Database PL/SQL Packages and Types Reference

Oracle Data Guard Concepts and Administration Guide

Question 3

Question Type: MultipleChoice

Examine this list of possible steps:

1. Raise the compatibility level on both databases.

2. Restart SQL Apply on the upgraded logical standby database.

3. Start SQL Apply on the old primary database.

4. Perform a Switchover to the logical standby database.

5. Upgrade the logical standby database.

6. Upgrade the old primary database.

Which is the minimum number of steps in the correct order, to perform a rolling release upgrade of a data guard environment using an existing logical standby database and to enable the new functionality?

Options:

A- 1,5,2,4,6,3

B- 5,2,4,6,3,1

C- 4,6,5,2,3,1

D- 5,2,4,1

E- 5,2,4,3,6,1

Answer:

A

Explanation:

The process of performing a rolling release upgrade in a Data Guard environment using a logical standby database generally involves these steps:

Raise the compatibility level on both databases (1): Ensuring both the primary and logical standby databases are operating with the same and correct compatibility level is essential before starting the upgrade process.

Upgrade the logical standby database (5): Apply the database upgrade to the logical standby first, which allows the primary database to continue serving the workload without interruption.

Restart SQL Apply on the upgraded logical standby database (2): Once the logical standby has been upgraded, SQL Apply must be restarted to apply the redo data from the primary database, which is still running the earlier version.

Perform a switchover to the logical standby database (4): After confirming that the logical standby database is successfully applying redo data, perform a switchover to make it the new primary database.

Upgrade the old primary database (6): With the new primary database now in place, upgrade the old primary database (which is now the new standby) to the new Oracle Database release.

Start SQL Apply on the old primary database (3): Finally, start SQL Apply on what is now the standby database to synchronize it with the new primary database. Reference:

Oracle Data Guard Concepts and Administration Guide

Oracle Database Upgrade Guide

Question 4

Question Type: **MultipleChoice**

Which three are prerequisites for enabling Fast-Start Failover?

Options:

A- The Data Guard environment must be managed by the Data Guard Broker.

B- Flashback Database must be enabled only on the Fast-Start Failover target standby database.

C- You can specify only one standby database as the fast-start failover target.

D- The configuration must be operating in either Maximum Performance or Maximum Protection mode.

E- The maximum protection mode can be used, but with two or more standby databases.

F- Flashback Database must be enabled on both the primary database and the Fast-Start Failover target standby database.

Answer:

A, C, F

Explanation:

To enable Fast-Start Failover in a Data Guard environment, the following conditions must be in place:

The Data Guard environment must be managed by the Data Guard Broker (A): The Broker simplifies management tasks and is required to enable fast-start failover, which is an automatic failover mechanism provided by Data Guard.

You can specify only one standby database as the fast-start failover target (C): Fast-start failover is designed to fail over to a single, predetermined standby database, known as the target standby.

Flashback Database must be enabled on both the primary database and the Fast-Start Failover target standby database (F): Flashback Database provides a quick way to revert a database to a point in time before a logical or physical corruption or error occurred. It must be enabled on both the primary and target standby databases to allow for the possibility of reinstating the old primary as a standby after a failover. Reference:

Oracle Data Guard Concepts and Administration Guide

Oracle Database High Availability Overview

Question 5

Question Type: MultipleChoice

Examine the Data Guard configuration:

DGMGRL> show configuration;

Configuration - Animals

Protection Mode: Max Availability

Databases:

dogs - Primary database sheep

- Physical standby database cats

- Physical standby database

Fast-Start Failover: DISABLED

Configuration Status: SUCCESS

An attempt to enable fast-start failover raises an error:

DGMGRL> enable fast_start failover;

Error: ORA-16693: requirements not met for enabling fast-start failover

Failed.

Identify three possible reasons for this error.

Options:

A- The fastStartFailoverTarget property is not set on Dogs.

B- The LogxptModr property is set to async on Sheep while Sheep is the target standby database.

C- The LogXptMode property is set to FASTSYNC on Cats while Sheep is the target standby database.

D- The LogXptMode property is set to async on Dogs.

E- The LogXptMode property is set to fastsync on Dogs.

Answer:

A, B, D

Explanation:

When enabling fast-start failover, certain conditions must be met:

The fastStartFailoverTarget property is not set on Dogs (A): The primary database (Dogs) needs to have a fast-start failover target configured for the operation to succeed.

The LogXptMode property is set to ASYNC on Sheep while Sheep is the target standby database (B): Fast-start failover requires synchronous redo transport (SYNC or FASTSYNC) to ensure zero data loss, which is a prerequisite for enabling the feature.

The LogXptMode property is set to ASYNC on Dogs (D): Similar to the previous point, the primary database must be configured to use synchronous redo transport for the fast-start failover to be possible. Reference:

Oracle Data Guard Broker documentation

Oracle Database Error Messages Guide

Question 6

Question Type: **MultipleChoice**

Your Data Guard environment contains a four-instance RAC primary database whose SID is PROD and a RAC physical standby database whose std is PROD_SBY.

Examine the command executed on a node of the primary database cluster to create a service OLTPWORKLOAD that the applications will use to connect to the database when it is in the FRIMARYTclatabase role:

srvctl add service -db PROD -service oltpworkload -role PRIMARY -failovertype SESSION -failovermethod BASIC -failoverdelay 10 -failoverretry 150

The service is then started

Consider this list of tasks:

1. On a node of the standby database cluster execute:

srvctl add service -db PROD_SBY -service oltpworkload -role PRIMARY -failovertype SESSION -failovermethod BASIC -failoverdelay 10 -failoverretry 150

2. On the primary database, create the oltpworkload database service using the dbms_service.create_service procedure.

3. Configure tap for clients in the tnsnames.ora files.

4. Make sure clients use the OLTPWORKLOAD service to connect to the database instances.

5. On the standby database, create the oltpworkload database service using the dbms_service.create_servi;l procedure.

Identify the required steps to configure and use Transparent Application Failover (taf).

Options:

A- 4

B- 2,3,4

C- 5

D- 1.4

E- 3,4

F- 1,3,4

Answer:

D

Explanation:

To set up Transparent Application Failover (TAF) in a Data Guard environment with RAC, you would need to:

On a node of the standby database cluster, execute the srvctl command to add the oltpworkload service for the PRIMARY role (1): This prepares the standby cluster to provide the oltpworkload service in case a failover occurs, and the standby becomes the primary database.

Make sure clients use the OLTPWORKLOAD service to connect to the database instances (4): This ensures that client connections are directed to the correct service, which is managed by TAF and can fail over in case of a primary database outage. Reference:

Oracle Real Application Clusters Administration and Deployment Guide

Oracle Data Guard Concepts and Administration Guide

Question 7

Question Type: MultipleChoice

Your Data Guard environment has two remote physical standby databases.

Client applications use the local naming method to connect to the primary database instance.

You want applications to automatically connect to the new primary database instance in case of a switchover or a failover.

Which set of actions will fulfill this requirement?

Options:

A- Set the LOCAL_LISTENER parameter for all the database instance to register services with the default listener on the primary database host.

B- Create a database service on the primary database that is started automatically by a trigger, when the database role is PRIMARY; modify the connection descriptors used by client applications to include all the standby hosts and connect to the database instance using that service name.

C- Set DB_NAME and DB_UNIQUE_NAME identically on all databases; modify the connection descriptors on client applications to include all the standby hosts and connect to the database instance using that service name.

D- Set the INSTANCE NAME parameter identically on all databases; modify the connection descriptor on client applications to include all the standby hosts and connect to the database instance using that service name.

Answer:

B

Explanation:

For seamless client redirection in a Data Guard environment, the following steps should be taken:

Create a database service on the primary database that is started automatically by a trigger when the database role is PRIMARY (B): This ensures that the service is only available on the primary database and is automatically started after a role transition due to switchover or failover.

Modify the connection descriptors used by client applications to include all the standby hosts and connect to the database instance using that service name (B): Client applications use the connection descriptors that include all potential primary hosts (i.e., the current primary and all standbys). This enables clients to connect to whichever database is currently acting as the primary using the service name.
Reference:

Oracle Data Guard Concepts and Administration Guide

Oracle Real Application Clusters Administration and Deployment Guide

Question 8

Question Type: MultipleChoice

Which TWO statements correctly describe the behavior of Automatic Block Media Recovery in a Data Guard environment, for a corrupt block in the example tablespace encountered by a session logged in as the SH user?

Options:

A- A corrupt block on the primary database can be automatically recovered, using a block from a standby database with Real-Time Query enabled.

B- A corrupt block on the primary database is automatically recovered, using a block from a flashback log from a standby database with

Real-Time Query enabled.

C- A corrupt block on a standby database with Real-Time Query enabled, is automatically recovered, using flashback logs from the standby database.

D- A corrupt block on a standby database with Real-Time Query enabled, can be automatically recovered, using a block from the primary database.

E- A corrupt block on the primary database is automatically recovered, using a block from a flashback log from the primary database.

Answer:

A, E

Explanation:

Automatic Block Media Recovery can be a significant feature for maintaining data integrity within a Data Guard configuration.

A corrupt block on the primary database can be automatically recovered, using a block from a standby database with Real-Time Query enabled (A): When a corrupted block is encountered on the primary database, Oracle can automatically replace it with a good block from the standby database where Real-Time Query is enabled, leveraging the standby as a source of good data.

A corrupt block on the primary database is automatically recovered, using a block from a flashback log from the primary database (E): If a good block version is available in the flashback logs of the primary database, Automatic Block Media Recovery can use it to recover the corrupted block on the primary. Reference:

Oracle Database Backup and Recovery User's Guide

Question 9

Question Type: **MultipleChoice**

Which TWO statements are true about Real-Time Query?

Options:

A- Setting standby_max_data_delay=0 requires synchronous redo transport.

B- Real-Time Query has no limitations regarding the protection level of the Data Guard environment.

C- Disabling Real-Time Query prevents the automatic start of redo apply when a physical standby databases opened read only.

D- Real-Time Query sessions can be connected to a Far Sync instance.

E- A standby database enabled for Real-Time Query cannot be the Fast-Start Failover target of the Data Guard configuration.

Answer:

A, C

Explanation:

Real-Time Query is a feature that allows queries to be run on a physical standby database while it is applying redo data. The relevant truths about it are:

Setting standby_max_data_delay=0 requires synchronous redo transport (A): For the real-time apply feature to function with no data delay (zero delay), synchronous redo transport must be used. This setting ensures that the data on the standby database is as current as possible before queries are executed against it.

Disabling Real-Time Query prevents the automatic start of redo apply when a physical standby database is opened read-only (C): If Real-Time Query is disabled, opening the standby database in read-only mode will not start the redo apply process automatically. Redo apply needs to be manually started to synchronize the standby database with the primary. Reference:

Oracle Data Guard Concepts and Administration Guide

Question 10

Question Type: MultipleChoice

Which THREE statements are true about Far Sync instances?

Options:

A- The Data Guard Broker must be used to deploy and manage Far Sync instances.

B- They work with any protection level.

C- They enable standby databases to be configured at remote distances from the primary without impacting performance on the primary.

D- They use an spfMe, a standby controlfile, and standby redo logs.

E- A primary database can ship redo directly to multiple Far Sync instances.

Answer:

A, C, E

Explanation:

Far Sync instances are a feature of Oracle Data Guard designed to support zero data loss protection over long distances:

The Data Guard Broker must be used to deploy and manage Far Sync instances (A): Data Guard Broker simplifies the deployment and management of Far Sync instances, which are an integral part of zero data loss protection configurations.

They enable standby databases to be configured at remote distances from the primary without impacting performance on the primary (C): Far Sync instances are designed to receive redo from the primary database and then forward it to a remote standby database, thereby avoiding any performance impact on the primary database itself.

A primary database can ship redo directly to multiple Far Sync instances (E): A primary database can be configured to send redo logs to more than one Far Sync instance, which can then forward the redo to their respective remote standby databases. Reference:

Oracle Data Guard Concepts and Administration Guide

Oracle Database High Availability Overview

Question 11

Question Type: MultipleChoice

Which TWO are benefits of using Transaction Guard in a Data Guard environment?

Options:

A- It protects against user errors being replicated to standby databases.

B- It provides application continuity by rolling back uncommitted transactions interrupted by a failover or switchover.

C- It protects against logical corruptions being replicated to standby databases.

D- It protects against recoverable errors during a planned or an unplanned outage of a primary database.

E- It provides application continuity by replaying transactions interrupted by a failover or a switchover

Answer:

B, D

Explanation:

Transaction Guard provides benefits in terms of transaction consistency and recovery in a Data Guard environment:

It provides application continuity by rolling back uncommitted transactions interrupted by a failover or switchover (B): Transaction Guard ensures that any uncommitted transactions at the time of an outage are rolled back consistently, thus preserving the integrity of the application's data and state.

It protects against recoverable errors during a planned or an unplanned outage of a primary database (D): Transaction Guard offers protection against errors that can occur during outages, allowing applications to resume operations more quickly and reliably after recovery. Reference:

Oracle Database High Availability Overview

Oracle Real Application Clusters Administration and Deployment Guide

Question 12

Question Type: MultipleChoice

Which two are true about managing and monitoring Oracle container databases in a Data Guard environment using the broker?

Options:

A- If the primary database is not a container database, then a standby may be a container database.

B- If the primary database is a container database, then a physical standby may be a non-container database.

C- If the primary database is a container database, then a logical standby may be a non-container database.

D- All broker actions execute at the root container for container databases.

E- After a role change, the broker opens all Pluggable databases (pdbb) on the new primary.

Answer:

D, E

Explanation:

In the context of Oracle Data Guard and container databases (CDBs) managed by Data Guard Broker:

All broker actions execute at the root container for container databases (D): When using Data Guard Broker to manage a CDB, the actions performed by the broker are executed at the level of the root container. This is because the root container maintains the control and configuration information that applies to the entire CDB, including all of its pluggable databases (PDBs).

After a role change, the broker opens all Pluggable databases (PDBs) on the new primary (E): Following a role transition such as a switchover or a failover, Data Guard Broker ensures that all PDBs within the CDB of the new primary database are opened, which is essential to resume operations of the PDBs without manual intervention. Reference:

Oracle Data Guard Broker documentation

Oracle Multitenant Administrator's Guide