DUMPSBOSS.

Oracle Database 19c: Data Guard Administration

Oracle 1z0-076

Version Demo

Total Demo Questions: 11

Total Premium Questions: 110

Buy Premium PDF

https://dumpsboss.co

support@dumpsboss.co

support@dumpsboss.co
dumpsboss.co

Topic Break Down

Topic	No. of Questions
Topic 1, Oracle Data Guard Architecture	22
Topic 2, Creating a Physical Standby Database	12
Topic 3, Creating a Logical Standby Database	12
Topic 4, Oracle Active Data Guard	20
Topic 5, Data Guard Protection Modes	10
Topic 6, Data Guard Broker	22
Topic 7, Data Guard Fast-Start Failover	6
Topic 8, Mix Questions	6
Total	110

QUESTION NO: 1

Examine the Data Guard configuration:

DGMGRL> show configuration;

Configuration - Animals

Protection Mode: MaxAvailability

Databases:

dogs - Primary database

cats - Snapshot standby database

sheep - Snapshot standby database

Fast-Start Failover: DISABLED

Configuration Status:

ORA-01034: ORACLE not available

ORA-16625: cannot reach database "dogs"

DGM-17017: unable to determine configuration status

ou wish to perform a failover to sheep. Which command, or sequence of commands, should you issue to the broker before executing failover to sheep; using the broker?

A. DGMGRL> convert database sheep to physical standby;

B. DGMGRL> convert database cats to physical standby;

C. DGMGRL> convert database sheep to physical standby;DGMGRL> edit configuration set protection mode as maxperformance;DGMGRL> convert database sheep to physical standby;

D. DGMGRL> convert database cats to physical standby;

E. None, because you can directly failover to a Snapshot Standby Database.

ANSWER: A

Explanation:

`DGMGRL> convert database sheep to physical standby;` is required before initiating the broker failover because a snapshot standby database cannot serve as a failover target. A snapshot standby is temporarily opened for read-write activity while preserving a guaranteed restore point, allowing its received redo to be retained but not applied. To make `sheep` eligible to assume the primary role, the broker must first convert it back to a physical standby database. During this conversion, the snapshot standby is reverted to the guaranteed restore point created when it entered snapshot mode, so any local changes made while it was open read-write are discarded. It then resumes its physical standby role and can apply available redo before the failover operation is issued. The unreachable primary, `dogs`, makes this an unplanned failover scenario, but it does not change the prerequisite that the intended target be a physical standby. The conversion is performed through the broker with the displayed `CONVERT DATABASE` command; after conversion completes, issue the appropriate broker `FAILOVER TO sheep` command. Oracle documents both the snapshot-to-physical standby conversion workflow and the requirement to use a physical standby as the failover target in its Data Guard Broker and standby database administration guides: Oracle Data Guard Broker documentation and Managing Physical Standby Databases.

QUESTION NO: 2

Which three are true about using Flashback database through role transitions in a Data Guard environment?

A. Physical standby databases retain their current role when you flash back to a point in time before a reinstate occurred which caused this database to become a physical standby.

B. Logical standby databases retain their current role when you flash back through to a point in time before the switchover occurred which caused this database to become a logical standby.

C. Logical standby database roles are reverted to their original role when you flash back to a point in time before the switchover occurred which caused this database to become a logical standby.

D. Physical standby databases retain their current role when you flash back to a point in time before the switchover occurred which caused this database to become a physical standby.

E. Flashback database may not be used to undo a physical standby database activation.

ANSWER: A C D

Explanation:

Physical standby databases retain their current role when they are flashed back across a switchover or reinstate operation. This behavior is important for Data Guard recovery workflows: Flashback Database restores the database contents to the requested earlier point while preserving the database's current physical-standby role, avoiding an unintended additional role transition. Thus, a database that became a physical standby as part of either a switchover or a Broker reinstate operation remains a physical standby after being flashed back before that operation.

Logical standby databases have distinct behavior when flashed back across the switchover that made them logical standbys. Their role is restored to the role that existed at the chosen flashback point. Therefore, flashing a logical standby to a time before that switchover reverts it to its original role. Administrators must account for this role behavior when planning flashback operations around logical-standby role transitions and must ensure the needed flashback logs are still available for the target time.

Oracle documents the Data Guard role-transition considerations for Flashback Database in the Oracle Database 19c Data Guard documentation. Related Flashback Database operation and prerequisites are described in the Oracle Database Backup and Recovery User's Guide.

QUESTION NO: 3

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?

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: E

Explanation:

SQL Tuning Advisor is the tool designed to analyze and tune SQL workloads executing on an Active Data Guard physical standby database. Active Data Guard permits the standby to remain open for read-only reporting while redo apply continues, so its reporting SQL can have execution plans, object-access patterns, and resource demands that differ significantly from those on the primary database. SQL Tuning Advisor analyzes individual SQL statements using the optimizer and relevant execution statistics, then produces actionable tuning recommendations. These can include SQL profiles, access-path or index-related recommendations, SQL restructuring guidance, and statistics-related recommendations, as applicable to the statement being analyzed.

Oracle Database 19c supports running SQL Tuning Advisor for SQL workload on an Active Data Guard standby, enabling tuning analysis to occur where the reporting workload actually runs and helping avoid unnecessary diagnostic load on the primary database. This makes it the appropriate purpose-built tool when the goal is SQL-level tuning rather than general workload history collection or broad performance diagnosis. Oracle documents SQL Tuning Advisor and its recommendation process in the Oracle Database SQL Tuning Guide; Active Data Guard capabilities are described in the Oracle Data Guard Concepts and Administration documentation.

QUESTION NO: 4

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?

A. The redo transport mode is set to ASYNC for the standby database Sheep.

B. The redo transport mode is set to ASYNC 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.

ANSWER: A

Explanation:

The redo transport mode is set to ASYNC for the standby database Sheep. Maximum Protection requires at least one physical standby database that can provide synchronous redo transport. Before a primary database transaction can commit under this protection mode, redo must be synchronously written to the online redo log of a qualifying standby. This requirement ensures that committed transactions are protected against data loss if the primary database fails.

In this configuration, Cats is a snapshot standby database and therefore cannot be the standby that supports Maximum Protection. Sheep is consequently the only physical standby that could satisfy the synchronous transport requirement. The ORA-16627 error means that, when the broker evaluates the configuration for the requested protection mode, no eligible standby remains to support it. With Sheep as the sole eligible physical standby, this establishes that its current redo transport mode is asynchronous rather than synchronous. The broker therefore refuses the protection-mode change rather than allowing a configuration that cannot meet Maximum Protection guarantees.

Oracle documents the requirements for protection modes and the need for synchronous transport in Maximum Protection in the Oracle Data Guard Protection Modes documentation. The broker's management of redo transport and protection mode is described in the Data Guard Broker documentation.

QUESTION NO: 5

A primary database and a physical standby database were created by using RMAN DUPLICATE. Redo transport does not start, and the primary alert log reports an error indicating that the Data Guard configuration identifiers do not match.

Both databases have `DB_NAME=PROD`, and both databases were configured with `DB_UNIQUE_NAME=PROD`.

What should be done to correct the configuration?

A. Set a different `DB_NAME` for the standby database and retain the same `DB_UNIQUE_NAME`.

B. Set a unique `DB_UNIQUE_NAME` for the standby database and update the Data Guard configuration settings accordingly.

C. Regenerate the standby control file so that its DBID differs from that of the primary database.

D. Configure the standby database with the same `ORACLE_SID` as the primary database.

ANSWER: B

Explanation:

The physical standby must have the same `DB_NAME` as the primary database, but it must have a different `DB_UNIQUE_NAME`. The `DB_UNIQUE_NAME` identifies an individual database within a Data Guard configuration and is used by redo transport and Data Guard configuration parameters.

Set a unique value, such as `PROD_SBY`, for the standby `DB_UNIQUE_NAME`, and ensure that the `LOG_ARCHIVE_CONFIG` and redo transport destination settings reference the correct unique database names. Changing `DB_NAME` would make the database unsuitable as a physical standby, while changing only the connect identifier does not correct the Data Guard identifier mismatch.

QUESTION NO: 6

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?

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: E

Explanation:

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. SQL Apply skip rules are managed through the `DBMS_LOGSTDBY` package. An `UNSKIP` call removes a previously registered rule only when its rule definition matches the corresponding `SKIP` call. This includes the statement type, schema pattern, object-name pattern, and—where applicable—the procedure name. The object patterns used when the filters were created identify `EMP_NEW` and `EMP_OLD` as separate rules. Supplying the broader `EMP%` pattern does not identify either stored rule exactly and does not act as a bulk removal operation. Therefore, SQL Apply raises an error rather than deleting both filters. Each filter must be removed with its own matching `UNSKIP` invocation, using the same object-name value or pattern used to create that specific filter. Oracle documents `UNSKIP` as removing a skip rule and specifies that its parameters must match those used for the associated `SKIP` rule. See the DBMS_LOGSTDBY package reference and Oracle's logical standby database administration guide.

QUESTION NO: 7

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?

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: B

Explanation:

The correct sequence is **5,2,4,6,3,1**. A logical standby database enables a rolling upgrade because it can be upgraded first while the original primary continues processing application work on the earlier release. After the logical standby is upgraded, SQL Apply is restarted so that it can resume applying redo received from the still-old-release primary. Once redo application has caught up sufficiently and the upgraded logical standby is ready, a switchover makes that upgraded database the primary.

The former primary, now operating as the logical standby, is then upgraded. SQL Apply is started on that upgraded former primary to restore normal logical standby redo-apply protection in the reversed configuration. Finally, the `COMPATIBLE` initialization parameter is raised on both databases. Keeping compatibility at the earlier value until both databases have been upgraded preserves the rolling-upgrade interoperability requirements; raising it only after the full transition enables use of functionality dependent on the new database release. Oracle documents this logical-standby rolling-upgrade workflow in the Oracle Data Guard Administration guide; related release-upgrade requirements are covered in the Oracle Database Upgrade Guide.

QUESTION NO: 8

Which THREE statements are true about Far Sync instances?

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 SPFILE, a far sync control file, and standby redo logs.

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

ANSWER: C D E

Explanation:

They enable standby databases to be configured at remote distances from the primary without impacting performance on the primary because a Far Sync instance is placed close to the primary database. The primary can synchronously transmit redo over the short local network connection to the Far Sync instance, while that instance forwards the redo asynchronously across the wide-area network to a remote standby database. This design provides the basis for zero-data-loss protection across long distances without making primary commit processing wait for WAN round-trip latency.

They use an SPFILE, a far sync control file, and standby redo logs because a Far Sync instance is an Oracle instance dedicated to receiving and forwarding redo. It does not host database data files or function as a recoverable standby database; however, it requires its parameter and control-file infrastructure and standby redo logs to receive synchronized redo and reliably forward it downstream.

A primary database can ship redo directly to multiple Far Sync instances. This permits a Data Guard design to use more than one Far Sync destination when required for availability, geographic routing, or protection of multiple remote standby configurations. Oracle documents Far Sync architecture and its redo-forwarding role in the Oracle Data Guard Far Sync documentation and the associated Far Sync creation documentation.

QUESTION NO: 9

Examine the following parameter settings of the physical standby database:

• STANDBY_FILE_MANAGEMENT=AUTO

• ENABLED_PDBS_ON_STANDBY=

During which TWO tasks are files automatically created in the physical standby database after structure changes on the primary database?

A. Performing transportable tablespaces

B. Adding or dropping a redo file group

C. Adding a data file or creating a tablespace

D. Creating a PDB from the existing PDB within the same CDB

E. Renaming a data file in the primary database

ANSWER: C D

Explanation:

Adding a data file or creating a tablespace is correct because `STANDBY_FILE_MANAGEMENT=AUTO` directs a physical standby to automatically maintain standby datafiles when redo apply processes structural changes made at the primary. When a primary tablespace is created or a datafile is added, the corresponding datafile is created on the standby, using the configured file-name conversion or Oracle Managed Files destination as applicable. This preserves the required physical consistency between the primary and standby databases without a manual standby-side datafile creation step.

Creating a PDB from the existing PDB within the same CDB is also correct. Creating a PDB produces PDB datafiles, and the associated redo is applied on the physical standby. With automatic standby file management enabled, the files needed for that PDB are created at the standby. A null `ENABLED_PDBS_ON_STANDBY` setting leaves PDB recovery enabled for the standby's PDBs by default, so the new PDB can be maintained through redo apply. Oracle documents automatic file management for Data Guard and the handling of PDB creation in a Data Guard configuration in the Oracle Data Guard physical standby administration guide and the Oracle Multitenant administration guide.

QUESTION NO: 10

Which TWO statements are true about configuring Oracle Net Service in a Data Guard environment?

A. A static service must be registered with the local listener to enable DGMGRL to restart instances during the course of broker operations.

B. Install the oracle-database-preinstall-19c package to set the kernel parameters for Oracle Net based on the Data Guard best practice guidelines.

C. Installing the oracle-database-preinstall-19c package is NOT sufficient to set up operating system kernel parameters for Oracle Net.

D. Enterprise Manager does not require static service registration to restart instances during the course of broker operations.

E. It is necessary to use the failover clause for an address_list with multiple address lists in the tnsnames.ora file.

ANSWER: A B

Explanation:

A static service must be registered with the local listener to enable DGMGRL to restart instances during the course of broker operations. The Data Guard broker must be able to establish an administrative connection to an instance when it is not open, or when dynamic service registration is unavailable during a startup or restart operation. A statically registered listener service supplies that connection path, allowing DGMGRL to perform broker-managed instance lifecycle operations reliably. Oracle documents this static listener requirement for broker instance management in its Data Guard Broker concepts documentation.

Install the oracle-database-preinstall-19c package to set the kernel parameters for Oracle Net based on the Data Guard best practice guidelines. On supported Linux distributions, the Oracle Database preinstallation RPM configures required operating-system prerequisites, including recommended kernel parameter settings and other installation dependencies. Using this package provides a consistent baseline for the Oracle Database software that supports Oracle Net connectivity and Data Guard deployments, rather than requiring every applicable operating-system prerequisite to be configured manually. Oracle describes the package's prerequisite and kernel-parameter configuration role in About Oracle Database Preinstallation RPM.

QUESTION NO: 11

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?

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 D

Explanation:

Automatic Block Media Recovery (ABMR), also called automatic block repair in Data Guard documentation, can obtain a known-good copy of an individual corrupted block from the database's Data Guard partner. When corruption is detected while reading a block on the primary, the primary can request the block from a physical standby database that is open for real-time query. The standby's real-time query state ensures that it is available to service the request while redo apply remains active, allowing Oracle to repair the primary block without requiring a manual block media recovery operation.

ABMR also works in the opposite direction for a physical standby database that is open with Real-Time Query. If a query on that standby encounters a corrupt block, Oracle can request a valid copy of the block from the primary database and use it to repair the standby copy. This bidirectional use of a primary and an Active Data Guard physical standby helps preserve availability for read workloads and minimizes the operational impact of isolated physical block corruption. The SH schema identifies the session encountering the block, but it does not alter the recovery mechanism; the relevant factors are the database role, the physical standby relationship, and Real-Time Query availability. Oracle documents this capability in its Data Guard Concepts and Administration guide and describes related corruption handling in the Database Administrator's Guide.