

Lead2Pass: Latest Free Oracle 1Z0-060 Dumps (41-50) Download!

QUESTION 41

You are planning the creation of a new multitenant container database (CDB) and want to store the ROOT and SEED container data files in separate directories.

You plan to create the database using SQL statements.

Which three techniques can you use to achieve this?

- A. Use Oracle Managed Files (OMF).
- B. Specify the SEED FILE_NAME_CONVERT clause.
- C. Specify the PDB_FILE_NAME_CONVERT initialization parameter.
- D. Specify the DB_FILE_NAME_CONVERT initialization parameter.
- E. Specify all files in the CREATE DATABASE statement without using Oracle managed Files (OMF).

Answer: ABC

Explanation:

You must specify the names and locations of the seed's files in one of the following ways:

- * (A) Oracle Managed Files
- * (B) The SEED FILE_NAME_CONVERT Clause
- * (C) The PDB_FILE_NAME_CONVERT Initialization Parameter

QUESTION 42

You are about to plug a multi-terabyte non-CDB into an existing multitenant container database (CDB).

The characteristics of the non-CDB are as follows:

- Version: Oracle Database 11gRelease 2 (11.2.0.2.0) 64-bit
- Character set: AL32UTF8
- National character set: AL16UTF16
- O/S: Oracle Linux 6 64-bit

The characteristics of the CDB are as follows:

- Version: Oracle Database 12c Release 1 64-bit
- Character Set: AL32UTF8
- National character set: AL16UTF16
- O/S: Oracle Linux 6 64-bit

Which technique should you use to minimize down time while plugging this non-CDB into the CDB?

- A. Transportable database
- B. Transportable tablespace
- C. Data Pump full export/import
- D. The DBMS_PDB package
- E. RMAN

Answer: D

Explanation:

*Overview, example:

- Log into ncdb12c as sys
- Get the database in a consistent state by shutting it down cleanly.
- Open the database in read only mode
- Run DBMS_PDB.DESCRIBE to create an XML file describing the database.
- Shut down ncdb12c
- Connect to target CDB (CDB2)

- Check whether non-cdb (NCDB12c) can be plugged into CDB(CDB2)
- Plug-in Non-CDB (NCDB12c) as PDB(NCDB12c) into target CDB(CDB2).
- Access the PDB and run the noncdb_to_pdb.sql script.
- Open the new PDB in read/write mode.

*You can easily plug an Oracle Database 12c non-CDB into a CDB. Just create a PDB manifest file for the non-CDB, and then use the manifest file to create a cloned PDB in the CDB. *Note that to plugin a non-CDB database into a CDB, the non-CDB database needs to be of version 12c as well. So existing 11g databases will need to be upgraded to 12c before they can be part of a 12c CDB.

QUESTION 43

Your database supports an online transaction processing (OLTP) application. The application is undergoing some major schema changes, such as addition of new indexes and materialized views. You want to check the impact of these changes on workload performance.

What should you use to achieve this?

- A. Database replay
- B. SQL Tuning Advisor
- C. SQL Access Advisor
- D. SQL Performance Analyzer
- E. Automatic Workload Repository compare reports

Answer: E

Explanation:

While an AWR report shows AWR data between two snapshots (or two points in time), the AWR Compare Periods report shows the difference between two periods (or two AWR reports with a total of four snapshots). Using the AWR Compare Periods report helps you to identify detailed performance attributes and configuration settings that differ between two time periods.

QUESTION 44

An administrator account is granted the CREATE SESSION and SET CONTAINER system privileges.

A multitenant container database (CDB) instant has the following parameter set:

THREADED_EXECUTION = FALSE

Which four statements are true about this administrator establishing connections to root in a CDB that has been opened in read only mode?

- A. You can conned as a common user by using the connect statement.
- B. You can connect as a local user by using the connect statement.
- C. You can connect by using easy connect.
- D. You can connect by using OS authentication.
- E. You can connect by using a Net Service name.
- F. You can connect as a local user by using the SET CONTAINER statement.

Answer: CDEF

Explanation:

*The choice of threading model is dictated by the THREADED_EXECUTION initialization parameter.

THREADED_EXECUTION=FALSE: The default value causes Oracle to run using the multiprocess model.

THREADED_EXECUTION=TRUE: Oracle runs with the multithreaded model. *OS Authentication is not supported with the multithreaded model.

*THREADED_EXECUTION

When this initialization parameter is set to TRUE, which enables the multithreaded Oracle model, operating system authentication is not supported. Attempts to connect to the database using operating system authentication (for example, CONNECT / AS SYSDBA or CONNECT /) when this initialization parameter is set to TRUE receive an ORA-01031"insufficient privileges" error.

F: The new SET CONTAINER statement within a call back function:

The advantage of SET CONTAINER is that the pool does not have to create a new connection to a PDB, if there is an existing connection to a different PDB. The pool can use the existing connection, and through SET CONTAINER, can connect to the desired PDB. This can be done using:

```
ALTER SESSION SET CONTAINER=<PDB Name>
```

This avoids the need to create a new connection from scratch.

QUESTION 45

Examine the following query output:

```
SQL> SELECT name, force_logging FROM v$database;
NAME          FORCE_LOGGING
-----
PROD          NO
```

You issue the following command to import tables into the hr schema:

```
$ > impdp hr/hr directory = dumpdir dumpfile = hr_new.dmp schemas=hr TRANSFORM=DISABLE_ARCHIVE_LOGGING: Y
```

Which statement is true?

- A. All database operations performed by the impdp command are logged.
- B. Only CREATE INDEX and CREATE TABLE statements generated by the import are logged.
- C. Only CREATE TABLE and ALTER TABLE statements generated by the import are logged.
- D. None of the operations against the master table used by Oracle Data Pump to coordinate its activities are logged.

Answer: D

Explanation:

* From the exhibit we see that FORCE_LOGGING is set to NO.

* Datapump Import impdp in 12c includes a new parameter to disable logging during data import. This option could improve performance of import tremendously during large data loads.

The TRANSFORM=DISABLE_ARCHIVE_LOGGING is used to disable logging. The value can be Y or N. Y to disable logging and N to enable logging.

However, if the database is running with FORCE LOGGING enabled, data pump ignores disable logging request.

Note:

* When the primary database is in FORCE LOGGING mode, all database data changes are logged. FORCE LOGGING mode ensures that the standby database remains consistent with the primary database.

* force_logging V\$database

A tablespace or the entire database is either in force logging or no force logging mode. To see which it is, run:

```
SQL> SELECT force_logging FROM v$database;
```

```
FOR
```

```
---
```

```
NO
```

QUESTION 46

You notice a performance change in your production Oracle database and you want to know which change has made this performance difference.

You generate the Compare Period Automatic Database Diagnostic Monitor (ADDM) report to further investigation.

Which three findings would you get from the report?

- A. It detects any configuration change that caused a performance difference in both time periods.

- B. It identifies any workload change that caused a performance difference in both time periods.
- C. It detects the top wait events causing performance degradation.
- D. It shows the resource usage for CPU, memory, and I/O in both time periods.
- E. It shows the difference in the size of memory pools in both time periods.
- F. It gives information about statistics collection in both time periods.

Answer: ABE

Explanation:

Keyword: shows the difference.

*Full ADDM analysis across two AWR snapshot periods

Detects causes, measure effects, then correlates them

Causes: workload changes, configuration changes

Effects: regressed SQL, reach resource limits (CPU, I/O, memory, interconnect) Makes actionable recommendations along with quantified impact *Identify what changed

/Configuration changes, workload changes

*Performance degradation of the database occurs when your database was performing optimally in the past, such as 6 months ago, but has gradually degraded to a point where it becomes noticeable to the users. The Automatic Workload Repository (AWR)

Compare Periods report enables you to compare database performance between two periods of time.

While an AWR report shows AWR data between two snapshots (or two points in time), the AWR Compare Periods report shows the difference (ABE) between two periods (or two AWR reports with a total of four snapshots). Using the AWR Compare Periods report helps you to identify detailed performance attributes and configuration settings that differ between two time periods.

QUESTION 47

Examine the parameter for your database instance:

NAME	TYPE	VALUE
optimizer_adaptive_reporting_only	boolean	FALSE
optimizer_capture_sql_plan_baselines	boolean	FALSE
optimizer_dynamic_sampling	integer	2
optimizer_features_enable	string	Lead2pass.com

You generated the execution plan for the following query in the plan table and noticed that the nested loop join was done. After actual execution of the query, you notice that the hash join was done in the execution plan:

Identify the reason why the optimizer chose different execution plans.

```
SQL> SELECT product_name
FROM   order_items o, product_information p
WHERE  o.unit_price = 15
AND    quantity > 1
AND    p.product_id = o.product_id;

30 rows selected.
```

- A. The optimizer used a dynamic plan for the query.
- B. The optimizer chose different plans because automatic dynamic sampling was enabled.
- C. The optimizer used re-optimization cardinality feedback for the query.
- D. The optimizer chose different plan because extended statistics were created for the columns used.

Answer: B

Explanation:

* optimizer_dynamic_sampling

OPTIMIZER_DYNAMIC_SAMPLING controls both when the database gathers dynamic statistics, and the size of the sample that the optimizer uses to gather the statistics.
Range of values 0 to 11

QUESTION 48

Which three statements are true about adaptive SQL plan management?

- A. It automatically performs verification or evolves non-accepted plans, in COMPREHENSIVE mode when they perform better than existing accepted plans.
- B. The optimizer always uses the fixed plan, if the fixed plan exists in the plan baseline.
- C. It adds new, better plans automatically as fixed plans to the baseline.
- D. The non-accepted plans are automatically accepted and become usable by the optimizer if they perform better than the existing accepted plans.
- E. The non-accepted plans in a SQL plan baseline are automatically evolved, in COMPREHENSIVE mode, during the nightly maintenance window and a persistent verification report is generated.

Answer: ADE

Explanation:

With adaptive SQL plan management, DBAs no longer have to manually run the verification or evolve process for non-accepted plans. When automatic SQL tuning is in COMPREHENSIVE mode, it runs a verification or evolve process for all SQL statements that have non-accepted plans during the nightly maintenance window. If the non-accepted plan performs better than the existing accepted plan (or plans) in the SQL plan baseline, then the plan is automatically accepted and becomes usable by the optimizer. After the verification is complete, a persistent report is generated detailing how the non-accepted plan performs compared to the accepted plan performance. Because the evolve process is now an AUTOTASK, DBAs can also schedule their own evolve job at end time.

Note:

*The optimizer is able to adapt plans on the fly by predetermining multiple subplans for portions of the plan.

*Adaptive plans, introduced in Oracle Database 12c, enable the optimizer to defer the final plan decision for a statement until execution time. The optimizer instruments its chosen plan (the default plan) with statistics collectors so that it can detect at runtime, if its cardinality estimates differ greatly from the actual number of rows seen by the operations in the plan. If there is a significant difference, then the plan or a portion of it will be automatically adapted to avoid suboptimal performance on the first execution of a SQL statement.

QUESTION 49

You create a new pluggable database, HR_PDB, from the seed database. Which three tablespaces are created by default in HR_PDB?

- A. SYSTEM
- B. SYSAUX
- C. EXAMPLE
- D. UNDO
- E. TEMP
- F. USERS

Answer: ABE

Explanation:

*A PDB would have its SYSTEM, SYSAUX, TEMP tablespaces. It can also contain other user-created tablespaces in it.

*

*Oracle Database creates both the SYSTEM and SYSAUX tablespaces as part of every database.

***tablespace_datafile_clauses**

Use these clauses to specify attributes for all data files comprising the SYSTEM and SYSAUX tablespaces in the seed PDB.

Incorrect:

Not D:a PDB can not have an undo tablespace. Instead, it uses the undo tablespace belonging to the CDB.

Note:

* Example:

```
CONN pdb_admin@pdb1
```

```
SELECT tablespace_name FROM dba_tablespaces;
```

```
TABLESPACE_NAME
```

```
-----
```

```
SYSTEM
```

```
SYSAUX
```

```
TEMP
```

```
USERS
```

```
SQL>
```

QUESTION 50

Which two statements are true about variable extent size support for large ASM files?

- A. The metadata used to track extents in SGA is reduced.
- B. Rebalance operations are completed faster than with a fixed extent size
- C. An ASM Instance automatically allocates an appropriate extent size.
- D. Resync operations are completed faster when a disk comes online after being taken offline.
- E. Performance improves in a stretch cluster configuration by reading from a local copy of an extent.

Answer: AC

Explanation:

A:Variable size extents enable support for larger ASM datafiles, reduce SGA memory requirements for very large databases(A), and improve performance for file create and open operations.

C:You don't have to worry about the sizes; the ASM instance automatically allocates the appropriate extent size.

Note:

*The contents of ASM files are stored in a disk group as a set, or collection, of data extents that are stored on individual disks within disk groups. Each extent resides on an individual disk. Extents consist of one or more allocation units (AU). To accommodate increasingly larger files, ASM uses variable size extents.

*The size of the extent map that defines a file can be smaller by a factor of 8 and 64 depending on the file size. The initial extent size is equal to the allocation unit size and it increases by a factor of 8 and 64 at predefined thresholds. This feature is automatic for newly created and resized datafiles when the disk group compatibility attributes are set to Oracle Release 11 or higher.

If you want to pass the Oracle 12c 1Z0-060 exam successfully, recommend to read latest Oracle 12c 1Z0-060 Dumps full version.

<http://www.lead2pass.com/1z0-060.html>