

PassLeader 1Z0-062 Exam Dumps Collection with VCE and PDF (Question 73 - Question 84)

PassLeader now are offering 100% pass ensure 1Z0-062 dumps! All 1Z0-062 exam questions have been updated with correct answers, welcome to download the newest PassLeader 1Z0-062 VCE dumps and PDF dumps:

<http://www.passleader.com/1z0-062.html> (210 Q&As --> 354 Q&As)

BTW: Download PassLeader 1Z0-062 dumps from Google Drive for free:

https://drive.google.com/open?id=0B-ob6L_QjGLpZVhq2dmaHZoZIE

QUESTION 73

Which two statements are true about Oracle Managed Files (OMF)?

- A. OMF cannot be used in a database that already has data files created with user-specified directions.
- B. The file system directions that are specified by OMF parameters are created automatically.
- C. OMF can be used with ASM disk groups, as well as with raw devices, for better file management.
- D. OMF automatically creates unique file names for table spaces and control files.
- E. OMF may affect the location of the redo log files and archived log files.

Answer: DE

Explanation:

D: The database internally uses standard file system interfaces to create and delete files as needed for the following database structures:

Tablespaces

Redo log files

Control files

Archived logs

Block change tracking files

Flashback logs

RMAN backups

Note:

* Using Oracle-managed files simplifies the administration of an Oracle Database. Oracle-managed files eliminate the need for you, the DBA, to directly manage the operating system files that make up an Oracle Database. With Oracle-managed files, you specify file system directories in which the database automatically creates, names, and manages files at the database object level. For example, you need only specify that you want to create a tablespace; you do not need to specify the name and path of the tablespace's datafile with the DATAFILE clause.

<http://www.oracle-base.com/articles/9i/oracle-managed-files.php> http://docs.oracle.com/cd/B10500_01/server.920/a96521/omf.htm

What Are Oracle-Managed Files?

QUESTION 74

Which four actions are possible during an Online Data file Move operation?

- A. Creating and dropping tables in the data file being moved
- B. Performing file shrink of the data file being moved
- C. Querying tables in the data file being moved
- D. Performing Block Media Recovery for a data block in the data file being moved
- E. Flashing back the database
- F. Executing DML statements on objects stored in the data file being moved

Answer: ACEF

Explanation:

You can now move On line Datafile without have to stop Monoged Recovery and manually copy and rename Files. This can even be used to move Datafiles from or to ASM.

- New in Oracle Database 12c: FROM METAUNK. Physical Standby Database is in Active Data Guard Mode (opened READ ONLY and Managed Recovery is running):

It is now possible to online move a Datafile while Managed Recovery is running, ie. the Physical Standby Database is in Active Data Guard Mode. You can use this Command to move the Datafile.

- A flashback operation does not relocate a moved data file to its previous location. If you move a data file online from one location to another and later flash back the database to a point in time before the move, then the Data file remains in the new location, but the contents of the Data file are changed to the contents at the time specified in the flashback.

QUESTION 75

Your multitenant container database (CDB) contains a pluggable database, HR_PDB. The default permanent tablespace in HR_PDB is USERDATA. The container database (CDB) is open and you connect RMAN. You want to issue the following RMAN command:

```
RMAN > BACKUP TABLESPACE hr_pdb:userdata;
```

Which task should you perform before issuing the command?

- A. Place the root container in ARCHIVELOG mode.
- B. Take the user data tablespace offline.
- C. Place the root container in the nomount stage.
- D. Ensure that HR_PDB is open.

Answer: A

QUESTION 76

Identify three scenarios in which you would recommend the use of SQL Performance Analyzer to analyze impact on the performance of SQL statements.

- A. Change in the Oracle Database version
- B. Change in your network infrastructure
- C. Change in the hardware configuration of the database server
- D. Migration of database storage from non-ASM to ASM storage
- E. Database and operating system upgrade

Answer: ACE

Explanation:

Oracle 11g/12c makes further use of SQL tuning sets with the SQL Performance Analyzer, which compares the performance of the statements in a tuning set before and after a database change. The database change can be as major or minor as you like, such as:

- * (E) Database, operating system, or hardware upgrades.
- * (A,C) Database, operating system, or hardware configuration changes.
- * Database initialization parameter changes.
- * Schema changes, such as adding indexes or materialized views.
- * Refreshing optimizer statistics.
- * Creating or changing SQL profiles.

QUESTION 77

Which two statements are true about the RMAN validate database command?

- A. It checks the database for intrablock corruptions.
- B. It can detect corrupt pfiles.
- C. It can detect corrupt spfiles.
- D. It checks the database for interblock corruptions.
- E. It can detect corrupt block change tracking files.

Answer: AC

QUESTION 78

You install a non-RAC Oracle Database. During Installation, the Oracle Universal Installer (OUI) prompts you to enter the path of the Inventory directory and also to specify an operating system group name. Which statement is true?

- A. The ORACLE_BASE base parameter is not set.
- B. The installation is being performed by the root user.
- C. The operating system group that is specified should have the root user as its member.
- D. The operating system group that is specified must have permission to write to the inventory directory.

Answer: D

Explanation:

Providing a UNIX Group Name

If you are installing a product on a UNIX system, the Installer will also prompt you to provide the name of the group which should own the base directory.

You must choose a UNIX group name which will have permissions to update, install, and deinstall Oracle software. Members of this group must have write permissions to the base directory chosen.

Only users who belong to this group are able to install or deinstall software on this machine.

QUESTION 79

You are required to migrate your 11.2.0.3 database as a pluggable database (PDB) to a multitenant container database (CDB). The following are the possible steps to accomplish this task:

1. Place all the user-defined tablespaces in read-only mode on the source database.
2. Upgrade the source database to a 12c version.
3. Create a new PDB in the target container database.
4. Perform a full transportable export on the source database with the VERSION parameter set to 12 using the expdp utility.
5. Copy the associated data files and export the dump file to the desired location in the target database.
6. Invoke the Data Pump import utility on the new PDB database as a user with the DATAPUMP_IMP_FULL_DATABASE role and specify the full transportable import options.
7. Synchronize the PDB on the target container database by using the DBMS_PDS.SYNC_ODB function.

Identify the correct order of the required steps.

- A. 2, 1, 3, 4, 5, 6
- B. 1, 3, 4, 5, 6, 7
- C. 1, 4, 3, 5, 6, 7
- D. 2, 1, 3, 4, 5, 6, 7
- E. 1, 5, 6, 4, 3, 2

Answer: C

Explanation:

1. Set user tablespaces in the source database to READ ONLY.
2. From the Oracle Database 11g Release 2 (11.2.0.3) environment, export the metadata and any data residing in administrative tablespaces from the source database using the FULL=Y and TRANSPORTABLE=ALWAYS parameters. Note that the VERSION=12 parameter is required only when exporting from an Oracle Database 11g Release 2 database.
3. Copy the tablespace data files from the source system to the destination system. Note that the log file from the export operation will list the data files required to be moved.
4. Create a COB on the destination system, including a PDB into which you will import the source database.
5. In the Oracle Database 12c environment, connect to the pre-created PDB and import the dump file. The act of importing the dump file will plug the tablespace data files into the destination PDB.

QUESTION 80

In your multitenant container database (CDB) with two pluggable database (PDBs). You want to create a new PDB by using SQL Developer. Which statement is true?

- A. The CDB must be open.
- B. The CDB must be in the mount stage.
- C. The CDB must be in the nomount stage.
- D. All existing PDBs must be closed.

Answer: A

Explanation:

* Creating a PDB

Rather than constructing the data dictionary tables that define an empty PDB from scratch, and then populating its Obj\$ and Dependency\$ tables, the empty PDB is created when the CDB is created. (Here, we use empty to mean containing no customer-created artifacts.) It is referred to as the seed PDB and has the name PDB\$Seed. Every CDB non-negotiably contains a seed PDB; it is non-negotiably always open in read-only mode. This has no conceptual significance; rather, it is just an optimization device. The create PDB operation is implemented as a special case of the clone PDB operation. The size of the seed PDB is only about 1 gigabyte and it takes only a few seconds on a typical machine to copy it.

QUESTION 81

Which two statements are true about the Oracle Direct Network File system (DNFS)?

- A. It utilizes the OS file system cache.
- B. A traditional NFS mount is not required when using Direct NFS.
- C. Oracle Disk Manager can manage NFS on its own, without using the operating kernel NFS driver.
- D. Direct NFS is available only in UNIX platforms.
- E. Direct NFS can load-balance I/O traffic across multiple network adapters.

Answer: CE

Explanation:

E: Performance is improved by load balancing across multiple network interfaces (if available).

Note:

* To enable Direct NFS Client, you must replace the standard Oracle Disk Manager (ODM) library with one that supports Direct NFS Client.

Incorrect:

Not A: Direct NFS Client is capable of performing concurrent direct I/O, which bypasses any operating system level caches and eliminates any operating system write-ordering locks.

Not B:

* To use Direct NFS Client, the NFS file systems must first be mounted and available over regular NFS mounts.

* Oracle Direct NFS (dNFS) is an optimized NFS (Network File System) client that provides faster and more scalable access to NFS storage located on NAS storage devices (accessible over TCP/IP).

Not D: Direct NFS is provided as part of the database kernel, and is thus available on all supported database platforms - even those that don't support NFS natively, like Windows.

Note:

* Oracle Direct NFS (dNFS) is an optimized NFS (Network File System) client that provides faster and more scalable access to NFS storage located on NAS storage devices (accessible over TCP/IP). Direct NFS is built directly into the database kernel - just like ASM which is mainly used when using DAS or SAN storage.

* Oracle Direct NFS (dNFS) is an internal I/O layer that provides faster access to large NFS files than traditional NFS clients.

QUESTION 82

Examine the parameters 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	12.1.0.1

Which three statements are true about the process of automatic optimization by using cardinality feedback?

- A. The optimizer automatically changes a plan during subsequent execution of a SQL statement if there is a huge difference in optimizer estimates and execution statistics.
- B. The optimizer can re optimize a query only once using cardinality feedback.
- C. The optimizer enables monitoring for cardinality feedback after the first execution of a query.
- D. The optimizer does not monitor cardinality feedback if dynamic sampling and multicolumn statistics are enabled.
- E. After the optimizer identifies a query as a re-optimization candidate, statistics collected by the collectors are submitted to the optimizer.

Answer: ACD

Explanation:

C: during the first execution of a SQL statement, an execution plan is generated as usual.

D: if multi-column statistics are not present for the relevant combination of columns, the optimizer can fall back on cardinality feedback.

(not B)* Cardinality feedback. This feature, enabled by default in 11.2, is intended to improve plans for repeated executions.

optimizer_dynamic_sampling

optimizer_features_enable

* dynamic sampling or multi-column statistics allow the optimizer to more accurately estimate selectivity of conjunctive predicates.

Note:

* OPTIMIZER_DYNAMIC_SAMPLING controls the level of dynamic sampling performed by the optimizer.

* Cardinality feedback was introduced in Oracle Database 11gR2. The purpose of this feature is to automatically improve plans for queries that are executed repeatedly, for which the optimizer does not estimate cardinalities in the plan properly. The optimizer may misestimate cardinalities for a variety of reasons, such as missing or inaccurate statistics, or complex predicates. Whatever the reason for the misestimate, cardinality feedback may be able to help.

QUESTION 83

Which three statements are true when the listener handles connection requests to an Oracle 12c database instance with multithreaded architecture enabled in UNIX?

- A. Thread creation must be routed through a dispatcher process.
- B. The local listener may spawn a new process and have that new process create a thread.
- C. Each Oracle process runs an SCM thread.
- D. Each multithreaded Oracle process has an SCM thread.
- E. The local listener may pass the request to an existing process which in turn will create a thread.

Answer: ADE

QUESTION 84

Which three operations can be performed as multipartition operations in Oracle?

- A. Merge partitions of a list partitioned table
- B. Drop partitions of a list partitioned table
- C. Coalesce partitions of a hash-partitioned global index
- D. Move partitions of a range-partitioned table
- E. Rename partitions of a range partitioned table
- F. Merge partitions of a reference partitioned index

Answer: ABF

Explanation:

Multipartition maintenance enables adding, dropping, truncate, merge, split operations on multiple partitions.

A: Merge Multiple Partitions:

The new "ALTER TABLE ... MERGE PARTITIONS" help merge multiple partitions or subpartitions with a single statement.

When merging multiple partitions, local and global index operations and semantics for inheritance of unspecified physical attributes are the same for merging two partitions.

B: Drop Multiple Partitions:

The new "ALTER TABLE ... DROP PARTITIONS" help drop multiple partitions or subpartitions with a single statement.

Example:

view plaincopy to clipboardprint?

```
SQL> ALTER TABLE Tab_tst1 DROP PARTITIONS
```

```
Tab_tst1_PART5, Tab_tst1_PART6, Tab_tst1_PART7;
```

Table altered

```
SQL>
```

Restrictions :

- You can't drop all partitions of the table.
- If the table has a single partition, you will get the error: ORA-14083: cannot drop the only partition of a partitioned.

PassLeader now are offering 100% pass ensure 1Z0-062 dumps! All 1Z0-062 exam questions have been updated with correct

answers, welcome to download the newest PassLeader 1Z0-062 VCE dumps and PDF dumps:

<http://www.passleader.com/1z0-062.html> (210 Q&As --> 354 Q&As)

BTW: Download PassLeader 1Z0-062 dumps from Google Drive for free:

https://drive.google.com/open?id=0B-ob6L_QjGLpZVhq2dmaHZoZIE