

## Valid 1Z0-062 Dumps with VCE and PDF for Free (Question 25 - Question 36)

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

### QUESTION 25

Which statement is true concerning dropping a pluggable database (PDB)?

- A. The PDB must be open in read-only mode.
- B. The PDB must be in mount state.
- C. The PDB must be unplugged.
- D. The PDB data files are always removed from disk.
- E. A dropped PDB can never be plugged back into a multitenant container database (CDB).

Answer: C

### QUESTION 26

You notice a high number of waits for the db file scattered read and db file sequential read events in the recent Automatic Database Diagnostic Monitor (ADDM) report. After further investigation, you find that queries are performing too many full table scans and indexes are not being used even though the filter columns are indexed. Identify three possible reasons for this.

- A. Missing or stale histogram statistics
- B. Undersized shared pool
- C. High clustering factor for the indexes
- D. High value for the DB\_FILE\_MULTIBLOCK\_READ\_COUNT parameter
- E. Oversized buffer cache

Answer: ACD

Explanation:

D: DB\_FILE\_MULTIBLOCK\_READ\_COUNT is one of the parameters you can use to minimize I/O during table scans. It specifies the maximum number of blocks read in one I/O operation during a sequential scan. The total number of I/Os needed to perform a full table scan depends on such factors as the size of the table, the multiblock read count, and whether parallel execution is being utilized for the operation.

### QUESTION 27

Which three features work together, to allow a SQL statement to have different cursors for the same statement based on different selectivity ranges?

- A. Bind Variable Peeking
- B. SQL Plan Baselines
- C. Adaptive Cursor Sharing
- D. Bind variable used in a SQL statement
- E. Literals in a SQL statement

Answer: ACE

Explanation:

\* In bind variable peeking (also known as bind peeking), the optimizer looks at the value in a bind variable when the database performs a hard parse of a statement.

When a query uses literals, the optimizer can use the literal values to find the best plan. However, when a query uses bind variables, the optimizer must select the best plan without the presence of literals in the SQL text. This task can be extremely difficult. By peeking at bind values the optimizer can determine the selectivity of a WHERE clause condition as if literals had been used, thereby improving the plan.

C: Oracle 11g/12g uses Adaptive Cursor Sharing to solve this problem by allowing the server to compare the effectiveness of execution plans between executions with different bind variable values. If it notices suboptimal plans, it allows certain bind variable

values, or ranges of values, to use alternate execution plans for the same statement. This functionality requires no additional configuration.

#### QUESTION 28

You notice a performance change in your production Oracle 12c database. You want to know which change caused this performance difference. Which method or feature should you use?

- A. Compare Period ADDM report
- B. AWR Compare Period report
- C. Active Session History (ASH) report
- D. Taking a new snapshot and comparing it with a preserved snapshot

Answer: B

Explanation:

The awrddrpt.sql report is the Automated Workload Repository Compare Period Report. The awrddrpt.sql script is located in the \$ORACLE\_HOME/rdbms/admin directory.

Incorrect:

Not A: Compare Period ADDM

Use this report to perform a high-level comparison of one workload replay to its capture or to another replay of the same capture. Only workload replays that contain at least 5 minutes of database time can be compared using this report.

#### QUESTION 29

You want to capture column group usage and gather extended statistics for better cardinality estimates for the CUSTOMERS table in the SH schema. Examine the following steps:

1. Issue the SELECT DBMS\_STATS.CREATE\_EXTENDED\_STATS ('SH', 'CUSTOMERS') FROM dual statement.
2. Execute the DBMS\_STATS.SEED\_COL\_USAGE (null, 'SH', 500) procedure.
3. Execute the required queries on the CUSTOMERS table.
4. Issue the SELECT DBMS\_STATS.REPORT\_COL\_USAGE ('SH', 'CUSTOMERS') FROM dual statement.

Identify the correct sequence of steps.

- A. 3, 2, 1, 4
- B. 2, 3, 4, 1
- C. 4, 1, 3, 2
- D. 3, 2, 4, 1

Answer: B

Explanation:

Step 1: (2). Seed column usage

Oracle must observe a representative workload, in order to determine the appropriate column groups. Using the new procedure DBMS\_STATS.SEED\_COL\_USAGE, you tell Oracle how long it should observe the workload.

Step 2: (3) You don't need to execute all of the queries in your work during this window. You can simply run explain plan for some of your longer running queries to ensure column group information is recorded for these queries.

Step 3: (1) Create the column groups

At this point you can get Oracle to automatically create the column groups for each of the tables based on the usage information captured during the monitoring window. You simply have to call the DBMS\_STATS.CREATE\_EXTENDED\_STATS function for each table. This function requires just two arguments, the schema name and the table name. From then on, statistics will be maintained for each column group whenever statistics are gathered on the table.

Note:

\* DBMS\_STATS.REPORT\_COL\_USAGE reports column usage information and records all the SQL operations the database has processed for a given object.

\* The Oracle SQL optimizer has always been ignorant of the implied relationships between data columns within the same table. While the optimizer has traditionally analyzed the distribution of values within a column, he does not collect value-based relationships between columns.

\* Creating extended statistics Here are the steps to create extended statistics for related table columns with dbms\_stats.created\_extended\_stats:

- 1 - The first step is to create column histograms for the related columns.

2 - Next, we run `dbms_stats.create_extended_stats` to relate the columns together. Unlike a traditional procedure that is invoked via an execute ("exec") statement, Oracle extended statistics are created via a select statement.

#### QUESTION 30

Which three statements are true about Automatic Workload Repository (AWR)?

- A. All AWR tables belong to the SYSTEM schema.
- B. The AWR data is stored in memory and in the database.
- C. The snapshots collected by AWR are used by the self-tuning components in the database.
- D. AWR computes time model statistics based on time usage for activities, which are displayed in the `v$SYS` time model and `V$SESS_TIME_MODEL` views.
- E. AWR contains system wide tracing and logging information.

Answer: BCE

Explanation:

\* A fundamental aspect of the workload repository is that it collects and persists database performance data in a manner that enables historical performance analysis. The mechanism for this is the AWR snapshot. On a periodic basis, AWR takes a "snapshot" of the current statistic values stored in the database instance's memory and persists them to its tables residing in the SYSAUX tablespace.

\* AWR is primarily designed to provide input to higherlevel components such as automatic tuning algorithms and advisors, but can also provide a wealth of information for the manual tuning process.

#### QUESTION 31

You upgraded your database from pre-12c to a multitenant container database (CDB) containing pluggable databases (PDBs).

Examine the query and its output:

```
SQL> SELECT * FROM v$PFILE_users;

```

USERNAME	SYSDB	SYSOP	SYSAS	SYSBA	SYSDG	SYSKM	CON_ID
SYSD	TRUE	TRUE	FALSE	FALSE	FALSE	FALSE	0

Which two tasks must you perform to add users with SYSBACKUP, SYSDG, and SYSKM privilege to the password file?

- A. Assign the appropriate operating system groups to SYSBACKUP, SYSDG, SYSKM.
- B. Grant SYSBACKUP, SYSDG, and SYSKM privileges to the intended users.
- C. Re-create the password file with SYSBACKUP, SYSDG, and SYSKM privilege and the FORCE argument set to No.
- D. Re-create the password file with SYSBACKUP, SYSDG, and SYSKM privilege, and FORCE arguments set to Yes.
- E. Re-create the password file in the Oracle Database 12c format.

Answer: BD

Explanation:

\* orapwd

/ You can create a database password file using the password file creation utility, ORAPWD.

The syntax of the ORAPWD command is as follows:

```
orapwd FILE=filename [ENTRIES=numusers] [FORCE={y/n}] [ASM={y/n}] [DBUNIQUENAME=dbname]
[FORMAT={ 12|legacy}] [SYSBACKUP={y/n}] [SYSDG={y/n}] [SYSKM={y/n}] [DELETE={y/n}] [INPUT_FILE=input-fname]
force - whether to overwrite existing file (optional),
```

\* v\$PFILE\_users

/ 12c: V\$PFILE\_USERS lists all users in the password file, and indicates whether the user has been granted the SYSDBA, SYSOPER, SYSASM, SYSBACKUP, SYSDG, and SYSKM privileges.

/ 10c: sts users who have been granted SYSDBA and SYSOPER privileges as derived from the password file.

ColumnDatatypeDescription

USERNAMEVARCHAR2(30)The name of the user that is contained in the password file SYSDBAVARCHAR2(5)If TRUE, the user can connect with SYSDBA privileges SYSOPERVERCHAR2(5)If TRUE, the user can connect with SYSOPER privileges

Incorrect:

Not E: The format of the v\$PFILE\_users file is already in 12c format.

#### QUESTION 32

An application accesses a small lookup table frequently. You notice that the required data blocks are getting aged out of the default buffer cache. How would you guarantee that the blocks for the table never age out?

- A. Configure the KEEP buffer pool and alter the table with the corresponding storage clause.

- B. Increase the database buffer cache size.
- C. Configure the RECYCLE buffer pool and alter the table with the corresponding storage clause.
- D. Configure Automatic Shared Memory Management.
- E. Configure Automatic Memory Management.

Answer: A

Explanation:

Schema objects are referenced with varying usage patterns; therefore, their cache behavior may be quite different. Multiple buffer pools enable you to address these differences. You can use a KEEP buffer pool to maintain objects in the buffer cache and a RECYCLE buffer pool to prevent objects from consuming unnecessary space in the cache. When an object is allocated to a cache, all blocks from that object are placed in that cache. Oracle maintains a DEFAULT buffer pool for objects that have not been assigned to one of the buffer pools.

#### QUESTION 33

You conned using SQL Plus to the root container of a multitenant container database (CDB) with SYSDBA privilege. The CDB has several pluggable databases (PDBs) open in the read/write mode. There are ongoing transactions in both the CDB and PDBs. What happens after issuing the SHUTDOWN TRANSACTIONAL statement?

- A. The shutdown proceeds immediately. The shutdown proceeds as soon as all transactions in the PDBs are either committed or rolled back.
- B. The shutdown proceeds as soon as all transactions in the CDB are either committed or rolled back.
- C. The shutdown proceeds as soon as all transactions in both the CDB and PDBs are either committed or rolled back.
- D. The statement results in an error because there are open PDBs.

Answer: B

#### QUESTION 34

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 35

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 11g Release 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: B

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 36

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: D

Explanation:

You can use the SQL Performance Analyzer to analyze the SQL performance impact of any type of system change. Examples of common system changes include:

Database upgrades

Configuration changes to the operating system, hardware, or database

Database initialization parameter changes

Schema changes, such as adding new indexes or materialized views

Gathering optimizer statistics

SQL tuning actions, such as creating SQL profiles

[http://docs.oracle.com/cd/B28359\\_01/server.111/b28318/intro.htm#CNCPT961](http://docs.oracle.com/cd/B28359_01/server.111/b28318/intro.htm#CNCPT961)

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