

## PassLeader Valid 1Z0-062 Dumps with VCE and PDF (Question 61 - Question 72)

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\\_QjGLpZVhqd2dmaHZoZlE](https://drive.google.com/open?id=0B-ob6L_QjGLpZVhqd2dmaHZoZlE)

### QUESTION 61

You execute the following PL/SQL:

```
BEGIN
DBMS_FGA.add_policy(
object_schema => 'JIM',
object_name => 'PRODUCTS',
policy_name => 'PROD_AUDIT',
audit_condition => 'PRICE > 10000',
audit_column => 'PRICE');
END;
/
```

[www.passleader.com](http://www.passleader.com)

Which two statements are true?

- A. Fine-Grained Auditing (FGA) is enabled for the PRICE column in the PRODUCTS table for SELECT statements only when a row with PRICE > 10000 is accessed.
- B. FGA is enabled for the PRODUCTS.PRICE column and an audit record is written whenever a row with PRICE > 10000 is accessed.
- C. FGA is enabled for all DML operations by JIM on the PRODUCTS.PRICE column.
- D. FGA is enabled for the PRICE column of the PRODUCTS table and the SQL statements is captured in the FGA audit trial.

Answer: AB

Explanation:

DBMS\_FGA.add\_policy

\* The DBMS\_FGA package provides fine-grained security functions

\* ADD\_POLICY Procedure

This procedure creates an audit policy using the supplied predicate as the audit condition.

Incorrect:

Not C: object\_schema

The schema of the object to be audited. (If NULL, the current log-on user schema is assumed.)

### QUESTION 62

You execute the following commands to audit database activities:

```
SQL > ALTER SYSTEM SET AUDIT_TRIAL=DB, EXTENDED SCOPE=SPFILE;
```

```
SQL > AUDIT SELECT TABLE, INSERT TABLE, DELETE TABLE BY JOHN By SESSION WHENEVER SUCCESSFUL;
```

Which statement is true about the audit record that generated when auditing after instance restarts?

- A. One audit record is created for every successful execution of a SELECT, INSERT OR DELETE command on a table, and contains the SQL text for the SQL Statements.
- B. One audit record is created for every successful execution of a SELECT, INSERT OR DELETE command, and contains the execution plan for the SQL statements.
- C. One audit record is created for the whole session if john successfully executes a SELECT, INSERT, or DELETE command, and contains the execution plan for the SQL statements.
- D. One audit record is created for the whole session if JOHN successfully executes a select command, and contains the SQL text and bind variables used.

E. One audit record is created for the whole session if john successfully executes a SELECT, INSERT, or DELETE command on a table, and contains the execution plan, SQL text, and bind variables used.

Answer: A

Explanation:

Note:

\* BY SESSION

In earlier releases, BY SESSION caused the database to write a single record for all SQL statements or operations of the same type executed on the same schema objects in the same session. Beginning with this release (11g) of Oracle Database, both BY SESSION and BY ACCESS cause Oracle Database to write one audit record for each audited statement and operation.

\* BY ACCESS

Specify BY ACCESS if you want Oracle Database to write one record for each audited statement and operation.

Note:

If you specify either a SQL statement shortcut or a system privilege that audits a data definition language (DDL) statement, then the database always audits by access. In all other cases, the database honors the BY SESSION or BY ACCESS specification.

\* For each audited operation, Oracle Database produces an audit record containing this information:

/ The user performing the operation

/ The type of operation

/ The object involved in the operation

/ The date and time of the operation

QUESTION 63

You support Oracle Database 12c Oracle Database 11g, and Oracle Database log on the same server. All databases of all versions use Automatic Storage Management (ASM). Which three statements are true about the ASM disk group compatibility attributes that are set for a disk group?

A. The ASM compatibility attribute controls the format of the disk group metadata.

B. RDBMS compatibility together with the database version determines whether a database Instance can mount the ASM disk group.

C. The RDBMS compatibility setting allows only databases set to the same version as the compatibility value, to mount the ASM disk group.

D. The ASM compatibility attribute determines some of the ASM features that may be used by the Oracle disk group.

E. The ADVN compatibility attribute determines the ACFS features that may be used by the Oracle 10 g database.

Answer: ABD

Explanation:

AD: The value for the disk group COMPATIBLE.ASM attribute determines the minimum software version for an Oracle ASM instance that can use the disk group. This setting also affects the format of the data structures for the Oracle ASM metadata on the disk.

B: The value for the disk group COMPATIBLE.RDBMS attribute determines the minimum COMPATIBLE database initialization parameter setting for any database instance that is allowed to use the disk group. Before advancing the COMPATIBLE.RDBMS attribute, ensure that the values for the COMPATIBLE initialization parameter for all of the databases that access the disk group are set to at least the value of the new setting for COMPATIBLE.RDBMS.

For example, if the COMPATIBLE initialization parameters of the databases are set to either 11.1 or 11.2, then COMPATIBLE.RDBMS can be set to any value between 10.1 and 11.1 inclusively.

Not E:

/The value for the disk group COMPATIBLE.ADVN attribute determines whether the disk group can contain Oracle ASM volumes. The value must be set to 11.2 or higher. Before setting this attribute, the COMPATIBLE.ASM value must be 11.2 or higher. Also, the Oracle ADVN volume drivers must be loaded in the supported environment.

/ You can create an Oracle ASM Dynamic Volume Manager (Oracle ADVN) volume in a disk group. The volume device associated with the dynamic volume can then be used to host an Oracle ACFS file system.

The compatibility parameters COMPATIBLE.ASM and COMPATIBLE.ADVN must be set to 11.2 or higher for the disk group.

Note:

\* The disk group attributes that determine compatibility are COMPATIBLE.ASM, COMPATIBLE.RDBMS, and

COMPATIBLE.ADV. The COMPATIBLE.ASM and COMPATIBLE.RDBMS attribute settings determine the minimum Oracle Database software version numbers that a system can use for Oracle ASM and the database instance types respectively. For example, if the Oracle ASM compatibility setting is 11.2, and RDBMS compatibility is set to 11.1, then the Oracle ASM software version must be at least 11.2, and the Oracle Database client software version must be at least 11.1. The COMPATIBLE.ADV attribute determines whether the Oracle ASM Dynamic Volume Manager feature can create a volume in a disk group.

#### QUESTION 64

To enable the Database Smart Flash Cache, you configure the following parameters:

```
DB_FLASH_CACHE_FILE = '/dev/flash_device_1', '/dev/flash_device_2'  
DB_FLASH_CACHE_SIZE=64G
```

What is the result when you start up the database instance?

- A. It results in an error because these parameter settings are invalid.
- B. One 64G flash cache file will be used.
- C. Two 64G flash cache files will be used.
- D. Two 32G flash cache files will be used.

Answer: A

#### QUESTION 65

You executed this command to create a password file:

```
$ orapwd file = orapworcl entries = 10 ignorecase = N
```

Which two statements are true about the password file?

- A. It will permit the use of uppercase passwords for database users who have been granted the SYSOPER role.
- B. It contains username and passwords of database users who are members of the OSOPER operating system group.
- C. It contains usernames and passwords of database users who are members of the OSDBA operating system group.
- D. It will permit the use of lowercase passwords for database users who have granted the SYSDBA role.
- E. It will not permit the use of mixed case passwords for the database users who have been granted the SYSDBA role.

Answer: AD

Explanation:

\* You can create a password file using the password file creation utility, ORAPWD.

\* Adding Users to a Password File

When you grant SYSDBA or SYSOPER privileges to a user, that user's name and privilege information are added to the password file. If the server does not have an EXCLUSIVE password file (that is, if the initialization parameter REMOTE\_LOGIN\_PASSWORDFILE is NONE or SHARED, or the password file is missing), Oracle Database issues an error if you attempt to grant these privileges.

A user's name remains in the password file only as long as that user has at least one of these two privileges. If you revoke both of these privileges, Oracle Database removes the user from the password file.

\* The syntax of the ORAPWD command is as follows:

```
ORAPWD FILE=filename [ENTRIES=numusers]  
[FORCE={Y|N}] [IGNORECASE={Y|N}] [NOSYSDBA={Y|N}]
```

\* IGNORECASE

If this argument is set to y, passwords are case-insensitive. That is, case is ignored when comparing the password that the user supplies during login with the password in the password file.

#### QUESTION 66

Identify three valid methods of opening, pluggable databases (PDBs).

- A. ALTER PLUGGABLE DATABASE OPEN ALL ISSUED from the root
- B. ALTER PLUGGABLE DATABASE OPEN ALL ISSUED from a PDB
- C. ALTER PLUGGABLE DATABASE PDB OPEN issued from the seed
- D. ALTER DATABASE PDB OPEN issued from the root
- E. ALTER DATABASE OPEN issued from that PDB
- F. ALTER PLUGGABLE DATABASE PDB OPEN issued from another PDB
- G. ALTER PLUGGABLE DATABASE OPEN issued from that PDB

Answer: AEG

Explanation:

E: You can perform all ALTER PLUGGABLE DATABASE tasks by connecting to a PDB and running the corresponding ALTER DATABASE statement. This functionality is provided to maintain backward compatibility for applications that have been migrated to a CDB environment.

AG: When you issue an ALTER PLUGGABLE DATABASE OPEN statement, READ WRITE is the default unless a PDB being opened belongs to a CDB that is used as a physical standby database, in which case READ ONLY is the default.

You can specify which PDBs to modify in the following ways:

List one or more PDBs.

Specify ALL to modify all of the PDBs.

Specify ALL EXCEPT to modify all of the PDBs, except for the PDBs listed.

QUESTION 67

You administer an online transaction processing (OLTP) system whose database is stored in Automatic Storage Management (ASM) and whose disk group use normal redundancy. One of the ASM disks goes offline, and is then dropped because it was not brought online before DISK\_REPAIR\_TIME elapsed. When the disk is replaced and added back to the disk group, the ensuing rebalance operation is too slow. Which two recommendations should you make to speed up the rebalance operation if this type of failure happens again?

- A. Increase the value of the ASM\_POWER\_LIMIT parameter.
- B. Set the DISK\_REPAIR\_TIME disk attribute to a lower value.
- C. Specify the statement that adds the disk back to the disk group.
- D. Increase the number of ASMB processes.
- E. Increase the number of DBWR\_IO\_SLAVES in the ASM instance.

Answer: AD

Explanation:

A: ASM\_POWER\_LIMIT specifies the maximum power on an Automatic Storage Management instance for disk rebalancing. The higher the limit, the faster rebalancing will complete. Lower values will take longer, but consume fewer processing and I/O resources.

D:

\* Normally a separate process is fired up to do that rebalance. This will take a certain amount of time. If you want it to happen faster, fire up more processes. You tell ASM it can add more processes by increasing the rebalance power.

\* ASMB.

ASM Background Process

Communicates with the ASM instance, managing storage and providing statistics.

Incorrect:

Not B: A higher, not a lower, value of DISK\_REPAIR\_TIME would be helpful here. Not E: If you implement database writer I/O slaves by setting the DBWR\_IO\_SLAVES parameter, you configure a single (master) DBWR process that has slave processes that are subservient to it. In addition, I/O slaves can be used to "simulate" asynchronous I/O on platforms that do not support asynchronous I/O or implement it inefficiently. Database I/O slaves provide non-blocking, asynchronous requests to simulate asynchronous I/O.

QUESTION 68

You are administering a database and you receive a requirement to apply the following restrictions:

1. A connection must be terminated after four unsuccessful login attempts by user.
2. A user should not be able to create more than four simultaneous sessions.
3. User session must be terminated after 15 minutes of inactivity.
4. Users must be prompted to change their passwords every 15 days.

How would you accomplish these requirements?

- A. by granting a secure application role to the users.
- B. by creating and assigning a profile to the users and setting the REMOTE\_OS\_AUTHENT parameter to FALSE.
- C. By creating and assigning a profile to the users and setting the SEC\_MAX\_FAILED\_LOGIN\_ATTEMPTS parameter to 4.
- D. By Implementing Fine-Grained Auditing (FGA) and setting the REMOTE\_LOGIN\_PASSWORD\_FILE parameter to NONE.
- E. By implementing the database resource Manager plan and setting the SEC\_MAX\_FAILED\_LOGIN\_ATTEMPTS parameters

to 4.

Answer: A

Explanation:

You can design your applications to automatically grant a role to the user who is trying to log in, provided the user meets criteria that you specify. To do so, you create a secure application role, which is a role that is associated with a PL/SQL procedure (or PL/SQL package that contains multiple procedures). The procedure validates the user: if the user fails the validation, then the user cannot log in. If the user passes the validation, then the procedure grants the user a role so that he or she can use the application. The user has this role only as long as he or she is logged in to the application. When the user logs out, the role is revoked.

Incorrect:

Not B: REMOTE\_OS\_AUTHENT specifies whether remote clients will be authenticated with the value of the OS\_AUTHENT\_PREFIX parameter.

Not C, Not E: SEC\_MAX\_FAILED\_LOGIN\_ATTEMPTS specifies the number of authentication attempts that can be made by a client on a connection to the server process. After the specified number of failure attempts, the connection will be automatically dropped by the server process.

Not D: REMOTE\_LOGIN\_PASSWORDFILE specifies whether Oracle checks for a password file.

Values:

shared: One or more databases can use the password file. The password file can contain SYS as well as non-SYS users.

exclusive: The password file can be used by only one database. The password file can contain SYS as well as non-SYS users.

none: Oracle ignores any password file. Therefore, privileged users must be authenticated by the operating system.

Note:

The REMOTE\_OS\_AUTHENT parameter is deprecated. It is retained for backward compatibility only.

QUESTION 69

A senior DBA asked you to execute the following command to improve performance:

```
SQL> ALTER TABLE subscribe log STORAGE (BUFFER_POOL recycle);
```

You checked the data in the SUBSCRIBE\_LOG table and found that it is a large table containing one million rows. What could be a reason for this recommendation?

- A. The keep pool is not configured.
- B. Automatic Workarea Management is not configured.
- C. Automatic Shared Memory Management is not enabled.
- D. The data blocks in the SUBSCRIBE\_LOG table are rarely accessed.
- E. All the queries on the SUBSCRIBE\_LOG table are rewritten to a materialized view.

Answer: D

Explanation:

The most of the rows in SUBSCRIBE\_LOG table are accessed once a week.

QUESTION 70

Which three tasks can be automatically performed by the Automatic Data Optimization feature of Information lifecycle Management (ILM)?

- A. Tracking the most recent read time for a table segment in a user tablespace
- B. Tracking the most recent write time for a table segment in a user tablespace
- C. Tracking insert time by row for table rows
- D. Tracking the most recent write time for a table block
- E. Tracking the most recent read time for a table segment in the SYSAUX tablespace
- F. Tracking the most recent write time for a table segment in the SYSAUX tablespace

Answer: ABD

Explanation:

Incorrect:

Not E, Not F: When Heat Map is enabled, all accesses are tracked by the in-memory activity tracking module. Objects in the SYSTEM and SYSAUX tablespaces are not tracked.

\* To implement your ILM strategy, you can use Heat Map in Oracle Database to track data access and modification.

Heat Map provides data access tracking at the segment-level and data modification tracking at the segment and row level.

\* To implement your ILM strategy, you can use Heat Map in Oracle Database to track data access and modification. You can also use Automatic Data Optimization (ADO) to automate the compression and movement of data between different tiers of storage within the database.

#### QUESTION 71

Which two partitioned table maintenance operations support asynchronous Global Index Maintenance in Oracle database 12c?

- A. ALTER TABLE SPLIT PARTITION
- B. ALTER TABLE MERGE PARTITION
- C. ALTER TABLE TRUNCATE PARTITION
- D. ALTER TABLE ADD PARTITION
- E. ALTER TABLE DROP PARTITION
- F. ALTER TABLE MOVE PARTITION

Answer: CE

Explanation:

Asynchronous Global Index Maintenance for DROP and TRUNCATE PARTITION This feature enables global index maintenance to be delayed and decoupled from a DROP and TRUNCATE partition without making a global index unusable. Enhancements include faster DROP and TRUNCATE partition operations and the ability to delay index maintenance to off-peak time.

#### QUESTION 72

You configure your database Instance to support shared server connections. Which two memory areas that are part of PGA are stored in SGA instead, for shared server connection?

- A. User session data
- B. Stack space
- C. Private SQL area
- D. Location of the runtime area for DML and DDL Statements
- E. Location of a part of the runtime area for SELECT statements

Answer: AC

Explanation:

A: PGA itself is subdivided. The UGA (User Global Area) contains session state information, including stuff like package-level variables, cursor state, etc. Note that, with shared server, the UGA is in the SGA. It has to be, because shared server means that the session state needs to be accessible to all server processes, as any one of them could be assigned a particular session. However, with dedicated server (which likely what you're using), the UGA is allocated in the PGA.

C: The Location of a private SQL area depends on the type of connection established for a session. If a session is connected through a dedicated server, private SQL areas are located in the server process' PGA. However, if a session is connected through a shared server, part of the private SQL area is kept in the SGA.

Note:

\* System global area (SGA)

The SGA is a group of shared memory structures, known as SGA components, that contain data and control information for one Oracle Database instance. The SGA is shared by all server and background processes. Examples of data stored in the SGA include cached data blocks and shared SQL areas.

\* Program global area (PGA)

A PGA is a memory region that contains data and control information for a server process. It is nonshared memory created by Oracle Database when a server process is started. Access to the PGA is exclusive to the server process. There is one PGA for each server process. Background processes also allocate their own PGAs. The total memory used by all individual PGAs is known as the total instance PGA memory, and the collection of individual PGAs is referred to as the total instance PGA, or just instance PGA. You use database initialization parameters to set the size of the instance PGA, not individual PGAs.

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)