



# 1Z0-060<sup>Q&As</sup>

Upgrade to Oracle Database 12c

**Pass Oracle 1Z0-060 Exam with 100% Guarantee**

Free Download Real Questions & Answers **PDF** and **VCE** file from:

<https://www.geekcert.com/1z0-060.html>

100% Passing Guarantee  
100% Money Back Assurance

Following Questions and Answers are all new published by Oracle  
Official Exam Center

-  **Instant Download** After Purchase
-  **100% Money Back** Guarantee
-  **365 Days** Free Update
-  **800,000+** Satisfied Customers





### QUESTION 1

An Automatic Database Diagnostic Monitor (ADDM) finding in your production database reports that the shared pool is inadequately sized. You diagnose that this is due to the different kinds of workloads and this occurs only during peak hours. The following are the parameter settings for the database instance: You want to balance the memory between the System Global Area (SGA) components depending on the workload.

NAME	TYPE	VALUE
lock_sga	boolean	FALSE
pre_page_sga	boolean	FALSE
sga_max_size	big integer	300M
sga_target	big integer	0
fast_start_mttr_target	integer	0
memory_max_target	big integer	0
memory_target	big integer	0
pga_aggregate_target	big integer	100M
sga_target	big integer	0

Which option would solve this problem?

- A. setting the PGA\_AGGREGATE\_TARGET parameter to 200M and the SGA\_MAX\_SIZE parameter to 400M
- B. setting the MEMORY\_TARGET and SGA\_MAX\_SIZE parameters to 400M
- C. setting the SGA\_TARGET parameter to 300M
- D. setting the SGA\_MAX\_SIZE parameter to 400M

Correct Answer: C

### QUESTION 2

Your database supports a DSS workload that involves the execution of complex queries: Currently, the library cache contains the ideal workload for analysis. You want to analyze some of the queries for an application that are cached in the library cache.

What must you do to receive recommendations about the efficient use of indexes and materialized views to improve query performance?

- A. Create a SQL Tuning Set (STS) that contains the queries cached in the library cache and run the SQL Tuning Advisor (STA) on the workload captured in the STS.
- B. Run the Automatic Workload Repository (AWR) report.
- C. Create an STS that contains the queries cached in the library cache and run the SQL Performance Analyzer (SPA) on the workload captured in the STS.
- D. Create an STS that contains the queries cached in the library cache and run the SQL Access Advisor on the workload captured in the STS.



E. Run the Automatic Database Diagnostic Monitor (ADDM).

Correct Answer: D

\*

SQL Access Advisor is primarily responsible for making schema modification recommendations, such as adding or dropping indexes and materialized views. SQL Tuning Advisor makes other types of recommendations, such as creating SQL profiles and restructuring SQL statements.

\*

The query optimizer can also help you tune SQL statements. By using SQL Tuning Advisor and SQL Access Advisor, you can invoke the query optimizer in advisory mode to examine a SQL statement or set of statements and determine how to improve their efficiency. SQL Tuning Advisor and SQL Access Advisor can make various recommendations, such as creating SQL profiles, restructuring SQL statements, creating additional indexes or materialized views, and refreshing optimizer statistics.

Note:

\*

Decision support system (DSS) workload

\*

The library cache is a shared pool memory structure that stores executable SQL and PL/SQL code. This cache contains the shared SQL and PL/SQL areas and control structures such as locks and library cache handles.

---

### QUESTION 3

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

Correct Answer: AC

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's 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.

---

#### QUESTION 4

You executed this command to create a password file:

```
$ orapwd file = orapwrdl 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.

Correct Answer: AD

\*

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 5

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.

Correct Answer: A

\* 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.

[Latest 1Z0-060 Dumps](#)

[1Z0-060 PDF Dumps](#)

[1Z0-060 Braindumps](#)