



1Z0-060^{Q&As}

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.pass4itsure.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**

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

You generate an execution plan for this query and the plan table shows a nested loop join in the plan/

```
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;
```

After query execution, you see that a hash join was performed: What explain this optimizer behavior?

- A. It used SQL plan baselines for the query.
- B. It chose different plans because extended statistics were created on the filter columns.
- C. It chose different plans because automatic dynamic sampling was enabled.
- D. It used an adaptive execution plan.

Correct Answer: C

QUESTION 2

A new report process containing a complex query is written, with high impact on the database. You want to collect basic statistics about query, such as the level of parallelism, total database time, and the number of I/O requests.

For the database instance STATISTICS_LEVEL, the initialization parameter is set to TYPICAL and the CONTROL_MANAGEMENT_PACK_ACCESS parameter is set to DIAGNOSTIC+TUNING.

What should you do to accomplish this task?

- A. Execute the query and view Active Session History (ASH) for information about the query.
- B. Enable SQL trace for the query.
- C. Create a database operation, execute the query, and use the DBMS_SQL_MONITOR.REPORT_SQL_MONITOR



function to view the report.

D. Use the DBMS_APPLICATION_INFO.SET_SESSION_LONGOPS procedure to monitor query execution and view the information from the V\$SESSION_LONGOPS view.

Correct Answer: C

The REPORT_SQL_MONITOR function is used to return a SQL monitoring report for a specific SQL statement.

Incorrect Answers:

A: Not interested in session statistics, only in statistics for the particular SQL query.

B: We are interested in statistics, not tracing.

D: SET_SESSION_LONGOPS Procedure This procedure sets a row in the V\$SESSION_LONGOPS view. This is a view that is used to indicate the on-going progress of a long running operation. Some Oracle functions, such as parallel execution and Server Managed Recovery, use rows in this view to indicate the status of, for example, a database backup.

Applications may use the SET_SESSION_LONGOPS procedure to advertise information on the progress of application specific long running tasks so that the progress can be monitored by way of the V\$SESSION_LONGOPS view.

QUESTION 3

Which three statements are true about the use of the query result cache? (Choose three.)

- A. Results obtained from a query rewrite to a materialized view may be stored in the query result cache.
- B. Stale results may still be obtained from the query result cache, if the query session's query_rewrite_integrity parameter is set to stale_tolerated.
- C. Results obtained from a join between a table and a view may be stored in the query result cache.
- D. Results from distributed queries are never stored in the query result cache.
- E. Results may not be obtained from the query result cache, when the query_rewrite_enabled parameter is set to false.
- F. Results from remote queries may sometimes be stored in the query result cache.

Correct Answer: ABE

QUESTION 4

You created an encrypted tablespace:



```
SQL> CREATE TABLESPACE securespace
      DATAFILE '/home/user/oradata/secure01.dbf'
      SIZE 150M
      ENCRYPTION USING '3DES168'
      DEFAULT STORAGE (ENCRYPT) ;
```

You then closed the encryption wallet because you were advised that this is secure.

Later in the day, you attempt to create the EMPLOYEES table in the SECURESPACE tablespace with the SALT option on the EMPLOYEE column.

Which is true about the result?

- A. It creates the table successfully but does not encrypt any inserted data in the EMPNAME column because the wallet must be opened to encrypt columns with SALT.
- B. It generates an error when creating the table because the wallet is closed.
- C. It creates the table successfully, and encrypts any inserted data in the EMPNAME column because the wallet needs to be open only for tablespace creation.
- D. It generates error when creating the table, because the salt option cannot be used with encrypted tablespaces.

Correct Answer: B

QUESTION 5

You are the DBA for an Oracle Database 11g Release 1 data warehouse database running on Oracle

Linux with little endianness.

You want to plug in a transportable tablespace with a 16K blocksize from an Oracle Database 10g

database that is running on a SPARC-based Oracle Solaris system with big endianness.

The data warehouse database has the following parameters set:

```
DB_BLOCK_SIZE=8192
DB_CACHE_SIZE=200M
DB_16K_CACHE_SIZE=0
SHARED_POOL_SIZE=400M
LARGE_POOL_SIZE=32M
JAVA_POOL_SIZE=0
PGA_AGGREGATE_TARGET=1200M
```

The following parameters are not explicitly set:



SGA_TARGET
SGA_MAX_SIZE
MEMORY_TARGET
MEMORY_MAX_TARGET

What action must you perform on the data warehouse database to prepare it for the transportable tablespace plug in operation?

- A. Do nothing, because Automatic Memory Management will size DB_16K_CACHE_SIZE automatically
- B. Set ALTER SYSTEM SET DB_CACHE_SIZE=192M followed by ALTER SYSTEM SET DB_16K_CACHE_SIZE=8M;
- C. Do nothing, because you cannot transport a tablespace between databases with different endian formats.
- D. Do nothing, because Automatic Shared Memory Management will size DB_16K_CACHE_SIZE automatically.
- E. Set ALTER SYSTEM SET DB_16K_CACHE_SIZE=8M and leave other parameters set to their current settings.

Correct Answer: B

QUESTION 6

In your Oracle 12c database, you create two roles: AUDITOR and AUDIT_ADMIN. You then execute the following commands:

```
SQL> GRANT create view, create table TO auditor;  
SQL> GRANT auditor TO audit_admin;
```

To create a privilege analysis policy for the AUDIT_ADMIN role, you execute the following PL/SQL block:

```
SQL>BEGIN  
    DBMS_PRIVILEGE_CAPTURE.CREATE_CAPTURE(  
        NAME           => 'AUDIT ROLES CAPTURE',  
        DESCRIPTION    => 'CAPTURES AUDIT ADMIN ROLE USE',  
        TYPE           => DBMS_PRIVILEGE_CAPTURE.G_ROLE,  
        ROLES          => ROLE_NAME_LIST('AUDIT ADMIN');  
        CONDITION      => 'SYS_CONTEXT(''USERENV'', ''MODULE'') = ''ACCOUNT PAYABLE'' )'  
    END;  
/
```

Which statement is correct about the procedure?

- A. It always captures the use of the CREATE VIEW and CREATE TABLE privileges.
- B. It only captures the use of the privileges used through the AUDIT_ADMIN role when the ACCOUNT PAYABLE module is used.



C. It always captures the use of the privileges exercised through the AUDITOR and AUDIT_ADMIN roles.

D. It captures the use of the privileges used through the AUDIT_ADMIN and AUDITOR roles when the ACCOUNT PAYABLE module is used.

Correct Answer: B

References: https://docs.oracle.com/database/121/DVADM/priv_analysis.htm#DVADM591

QUESTION 7

Which two statements are true about Oracle Direct Network File System (NFS)?

A. Direct NFS can use only the orafnstab configuration file located in \$ORACLE_HOME/dbs/.

B. It utilizes the operating system file system cache.

C. Oracle Database requires mount points to be mounted by the kernel NFS system even when served through the Direct NFS client.

D. The Oracle database files available on the NFS server that are served by the Direct NFS Client cannot be accessed through the operating system kernel NFS.

E. Direct NFS performs load balancing across all specified paths. If a specified path fails, Direct NFS reissues I/O commands over any remaining paths.

Correct Answer: CE

QUESTION 8

Which three statements are true about SQL plan directives?

A. They are tied to a specific statement or SQL ID.

B. They instruct the maintenance job to collect missing statistics or perform dynamic sampling to generate a more optimal plan.

C. They are used to gather only missing statistics.

D. They are created for a query expression where statistics are missing or the cardinality estimates by the optimizer are incorrect.

E. They instruct the optimizer to create only column group statistics.

F. Improve plan accuracy by persisting both compilation and execution statistics in the SYSAUX tablespace.

Correct Answer: BDF

QUESTION 9

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	12.1.0.1

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:

```
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.
```

Identify the reason why the optimizer chose different execution plans.

- 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.

Correct Answer: B

* 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 10

While investigating a possible security breach, you realized that some audit records were not retrieved when the UNIFIED_AUDIT_TRAIL view was queried.

How can you ensure that all available audit records are returned when the audit trail is queried?

- A. Query the v\$UNIFIED_AUDIT_TRAIL view rather than the UNIFIED_AUDIT_TRAIL view.
- B. Enable the OS auditing option in the database before querying the UNIFIED_AUDIT_TRAIL view.
- C. Configure the IMMEDIATE write mode using the DBMS_AUDIT_MGMT.SET_AUDIT_TRAIL view.
- D. Execute the DBMS_AUDIT_MGMT.FLUSH_UNIFIED_AUDIT_TRAIL procedure before querying the UNIFIED_AUDIT_TRAIL view.



Correct Answer: D

Reference: <https://blogs.oracle.com/oraclemagazine/unify-auditing>

QUESTION 11

You are administering a database stored in Automatic Storage management (ASM). The files are stored in the DATA disk group. You execute the following command:

```
SQL > ALTER DISKGROUP data ADD ALIAS '+data/prod/myfile.dbf' FOR '+data.231.45678'
```

What is the result?

- A. The file '+data.231.45678' is physically relocated to '+data/prod' and renamed as 'myfile.dbf'.
- B. The file '+data.231.45678' is renamed as 'myfile.dbf', and copied to '+data/prod'.
- C. The file '+data.231.45678' remains in the same location and a synonym 'myfile.dbf' is created.
- D. The file 'myfile.dbf' is created in '+data/prod' and the reference to '+data.231.45678' in the data dictionary removed.

Correct Answer: C

ADD ALIAS Use this clause to create an alias name for an Oracle ASM filename. The alias_name consists of the full directory path and the alias itself.

QUESTION 12

Which three statements are true concerning unplugging a pluggable database (PDB)? (Choose three.)

- A. The unplugged PDB can be plugged into the same multitenant container database (CDB).
- B. The unplugged PDB can be plugged into another CDB.
- C. The PDB must be open in read-only mode.
- D. The PDB data files are automatically removed from disk.
- E. The unplugged PDB becomes a non-CDB.
- F. The PDB must be closed.

Correct Answer: ABF

Reference: http://www.oracle.com/webfolder/technetwork/tutorials/obe/db/12c/r1/pdb/pdb_unplug_plug/pdb_unplug_plug.html

QUESTION 13



Which three can be performed on a Data Pump type external table? (Choose three.)

- A. creating a public synonym for the table
- B. creating a view on the table
- C. partitioning the table
- D. performing a CREATE TABLE AS SELECT (CTAS) to populate the table
- E. creating an invisible index on the table
- F. updating the table using an UPDATE statement

Correct Answer: DEF

QUESTION 14

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.

Correct Answer: C

Note from Oracle Documentation: With redo logging disabled, the disk space required for redo logs during an Oracle Data Pump import will be smaller. However, to ensure recovery from media failure, the DBA should do an RMAN backup after the import completes. Even with this parameter specified, there is still redo logging for other operations of Oracle Data Pump. This includes all CREATE and ALTER statements, except CREATE INDEX, and all operations against the master table used by Oracle Data Pump during the import.

QUESTION 15



A database is stored in an Automatic Storage Management (ASM) disk group, disk group, DGROUP1 with SQL:

```
SQL> CREATE DISKGROUP dgroup1 NORMAL REDUNDANCY  
      FAILGROUP controller1 DISK '/devices/diska1', '/devices/diska2'  
      FAILGROUP controller2 DISK '/devices/diskb1', 'devices/diskb2';
```

There is enough free space in the disk group for mirroring to be done.

What happens if the CONTROLLER1 failure group becomes unavailable due to error or for maintenance?

- A. Transactions and queries accessing database objects contained in any tablespace stored in DGROUP1 will fail.
- B. Mirroring of allocation units will be done to ASM disks in the CONTROLLER2 failure group until the CONTROLLER1 failure group is brought back online.
- C. The data in the CONTROLLER1 failure group is copied to the CONTROLLER2 failure group and rebalancing is initiated.
- D. ASM does not mirror any data until the CONTROLLER1 failure group is brought back online, and newly allocated primary allocation units (AU) are stored in the CONTROLLER2 failure group, without mirroring.

Correct Answer: D

[Latest 1Z0-060 Dumps](#)

[1Z0-060 PDF Dumps](#)

[1Z0-060 VCE Dumps](#)