

**100%** Money Back  
**Guarantee**

**Vendor:**Oracle

**Exam Code:**1Z0-060

**Exam Name:**Upgrade to Oracle Database 12c

**Version:**Demo

## QUESTION 1

Your database has the SRV1 service configured for an application that runs on middle-tier application server. The application has multiple modules. You enable tracing at the service level by executing the following command:

```
SQL > exec DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE ('SRV1');
```

The possible outcomes of this command and the actions to aggregate trace files are as follows:

1.

The command fails because a module name is not specified.

2.

A trace file is created for each session that is running the SRV1 service.

3.

An aggregated trace file is created for all the sessions that are running the SRV1 service.

4.

Trace files may be aggregated by using the trcess utility.

5.

Trace files may be aggregated by using the tkprof utility.

Identify the correct outcome and the action to aggregate the trace files.

A. 1

B. 2 and 4

C. 2 and 5

D. 3 and 4

E. 3 and 5

Correct Answer: B

Tracing information is present in multiple trace files and you must use the trcess tool to collect it into a single file.

Note:

\*

The procedure enables a trace for a given combination of Service, MODULE and ACTION name. The specification is strictly hierarchical: Service Name or Service Name/MODULE, or Service Name, MODULE, and ACTION name must be specified. Omitting a qualifier behaves like a wild-card, so that not specifying an ACTION means all ACTIONS. Using the ALL\_ACTIONS constant achieves the same purpose.

\*

#### SERV\_MOD\_ACT\_TRACE\_ENABLE Procedure

This procedure will enable SQL tracing for a given combination of Service Name, MODULE and ACTION globally unless an instance\_name is specified.

\*

```
DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE(  
service_name IN VARCHAR2,  
module_name IN VARCHAR2 DEFAULT ANY_MODULE,  
action_name IN VARCHAR2 DEFAULT ANY_ACTION,  
waits IN BOOLEAN DEFAULT TRUE,  
binds IN BOOLEAN DEFAULT FALSE,  
instance_name IN VARCHAR2 DEFAULT NULL);
```

---

#### QUESTION 2

You execute this piece of code with appropriate privileges:

```
BEGIN  
  DBMS_REDACT.ADD_POLICY(  
    OBJECT_SCHEMA => 'SCOTT',  
    OBJECT_NAME   => 'EMP',  
    POLICY_NAME   => 'SCOTT_EMP'  
    COLUMN_NAME   => 'SAL',  
    FUNCTION_TYPE => DBMS_REDACT.FULL,  
  );  
END;  
/  
  
CREATE VIEW SCOTT.EMP_V AS SELECT * FROM SCOTT.EMP;  
  
BEGIN  
  DBMS_REDACT.ADD_POLICY(  
    OBJECT_SCHEMA => 'SCOTT',  
    OBJECT_NAME   => 'EMP_V',  
    POLICY_NAME   => 'SCOTT_EMP_V'  
    COLUMN_NAME   => 'SAL',  
    FUNCTION_TYPE => DBMS_REDACT.NONE,  
    EXPRESSION    => 'SYS_CONTEXT(''SYS_SESSION_ROLES'', ''MGR'') = ''FALSE''';  
  );  
END;  
/
```

User SCOTT has been granted the CREATE SESSION privilege and the MGR role.

Which two statements are true when a session logged in as SCOTT queries the SAL columns in the view and the table?

- A. Data is redacted for the EMP.SAL column only if the SCOTT session does not have the MGR role set.
- B. Data is never redacted on the EMP.SAL column.
- C. Data is redacted for the EMP.SAL column only if the SCOTT session has the MGR role set.
- D. Data is redacted for the EMP\_V.SAL column only if the SCOTT session has the MGR role set.
- E. Data is never redacted on the EMP\_V.SAL column.

Correct Answer: AE

---

### QUESTION 3

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

Correct Answer: D

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

References: [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)

---

### QUESTION 4

The tnsnames.ora file has an entry for the service alias ORCL as follows:

```

ORCL =
  ( DESCRIPTION =
    (ADDRESS = (PROTOCOL = TCP) (HOST = 10.156.24.216) (PORT = 1521))
    (CONNECT_DATA =
      (SERVER = DEDICATED)
      (SERVICE_NAME = orcl.oracle.com)
    )
  )

```

The TNSPING command executes successfully when tested with ORCL; however, from the same OS user session, you are not able to connect to the database instance with the following command:

```
SQL > CONNECT scott/tiger@orcl
```

What could be the reason for this?

- A. The listener is not running on the database node.
- B. The TNS\_ADMIN environment variable is set to the wrong value.
- C. The orcl.oracle.com database service is not registered with the listener.
- D. The DEFAULT\_DOMAIN parameter is set to the wrong value in the sqlnet.ora file.
- E. The listener is running on a different port.

Correct Answer: C

Service registration enables the listener to determine whether a database service and its service handlers are available. A service handler is a dedicated server process or dispatcher that acts as a connection point to a database. During registration, the LREG process provides the listener with the instance name, database service names, and the type and addresses of service handlers. This information enables the listener to start a service handler when a client request arrives.

## QUESTION 5

Your multitenant container database (CDB) contains three pluggable database (PDBs). You find that the control file is damaged. You plan to use RMAN to recover the control file. There are no startup triggers associated with the PDBs.

Which three steps should you perform to recover the control file and make the database fully operational? (Choose three.)

- A. Mount the container database (CDB) and restore the control file from the control file autobackup.
- B. Recover and open the CDB in NORMAL mode.
- C. Mount the CDB and then recover and open the database, with the RESETLOGS option.
- D. Open all the pluggable databases.
- E. Recover each pluggable database.
- F. Start the database instance in the nomount stage and restore the control file from control file autobackup.

Correct Answer: CDF

Step 1: Start the database instance in the nomount stage and restore the control file from control file auto backup

Step 2: Open all the pluggable databases.

Step 3: If all copies of the current control file are lost or damaged, then you must restore and mount a backup control file. You must then run the RECOVER command, even if no data files have been restored, and open the database with the RESETLOGS option.

Note:

\* RMAN and Oracle Enterprise Manager Cloud Control (Cloud Control) provide full support for backup and recovery in a multitenant environment. You can back up and recover a whole multitenant container database (CDB), root only, or one or more pluggable databases (PDBs).

---

## QUESTION 6

Your multitenant container database (CDB) is running in ARCHIVELOG mode. You connect to the CDB RMAN.

Examine the following command and its output:

```
RMAN> SELECT con_id, name, open_mode FROM V$PDBS;
  CON_ID NAME                                OPEN_MODE
-----
      2 PDB$SEED                                READ ONLY
      3 PDB2_1                                  MOUNTED
      4 PDB2_2                                  MOUNTED
```

You execute the following command:

```
RMAN > BACKUP DATABASE PLUS ARCHIVELOG;
```

Which data files will be backed up?

- A. Data files that belong to only the root container
- B. Data files that belong to the root container and all the pluggable databases (PDBs)
- C. Data files that belong to only the root container and PDB\$SEED
- D. Data files that belong to the root container and all the PDBs excluding PDB\$SEED

Correct Answer: B

Backing Up a Whole CDB

Backing up a whole CDB is similar to backing up a non-CDB. When you back up a whole CDB, RMAN backs up the root, all the PDBs, and the archived redo logs. You can then recover either the whole CDB, the root only, or one or more PDBs from the CDB backup.

Note:

\*

You can back up and recover a whole CDB, the root only, or one or more PDBs.

\*

### Backing Up Archived Redo Logs with RMAN

Archived redo logs are the key to successful media recovery. Back them up regularly. You can back up logs with `BACKUP ARCHIVELOG`, or back up logs while backing up datafiles and control files by specifying `BACKUP ... PLUS ARCHIVELOG`.

---

## QUESTION 7

What is the effect of specifying the "ENABLE PLUGGABLE DATABASE" clause in a "CREATE DATABASE" statement?

- A. It will create a multitenant container database (CDB) with only the root opened.
- B. It will create a CDB with root opened and seed read only.
- C. It will create a CDB with root and seed opened and one PDB mounted.
- D. It will create a CDB that must be plugged into an existing CDB.
- E. It will create a CDB with root opened and seed mounted.

Correct Answer: B

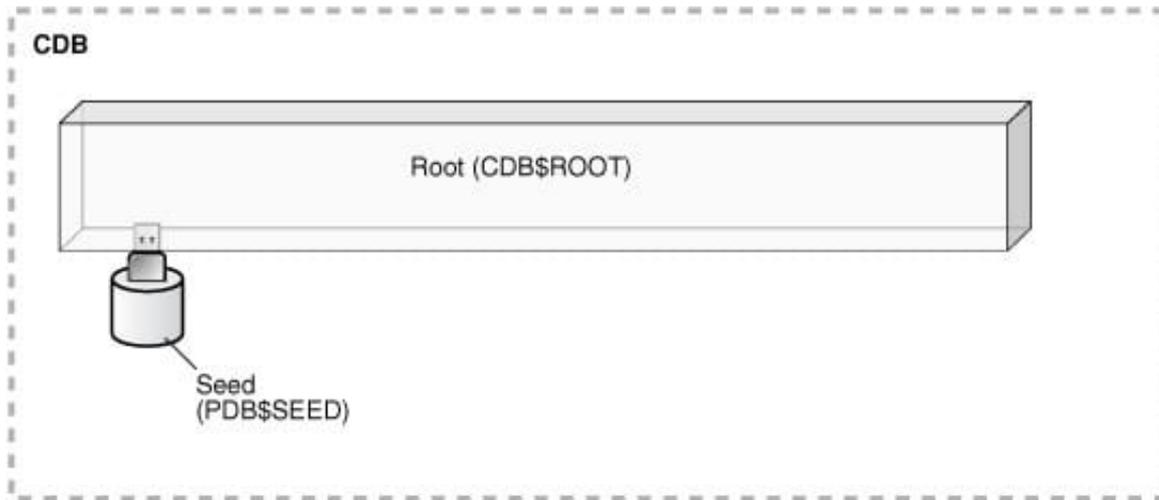
\*

The `CREATE DATABASE ... ENABLE PLUGGABLE DATABASE` SQL statement creates a new CDB. If you do not specify the `ENABLE PLUGGABLE DATABASE` clause, then the newly created database is a non-CDB and can never contain PDBs.

Along with the root (CDB\$ROOT), Oracle Database automatically creates a seed PDB (PDB\$SEED). The following graphic shows a newly created CDB:

\*

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.



### QUESTION 8

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.

Correct Answer: A

### QUESTION 9

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 10

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

Correct Answer: ACD

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 11

Examine the following steps of privilege analysis for checking and revoking excessive, unused privileges granted to users:

1.

Create a policy to capture the privilege used by a user for privilege analysis.

2.

Generate a report with the data captured for a specified privilege capture.

3.

Start analyzing the data captured by the policy.

4.

Revoke the unused privileges.

5.

Compare the used and unused privileges\ lists.

6.

Stop analyzing the data.

Identify the correct sequence of steps.

A. 1, 3, 5, 6, 2, 4

B. 1, 3, 6, 2, 5, 4

C. 1, 3, 2, 5, 6, 4

D. 1, 3, 2, 5, 6, 4

E. 1, 3, 5, 2, 6, 4

Correct Answer: B

1. Create a policy to capture the privilege used by a user for privilege analysis.

3. Start analyzing the data captured by the policy.

6. Stop analyzing the data.

2. Generate a report with the data captured for a specified privilege capture.

5. Compare the used and unused privileges\ lists.

4. Revoke the unused privileges.

---

## QUESTION 12

DBS\_IND\_PARTITIONS.ORPHANED\_ENTRIES contains YES for several index partitions. This is due to the use of deferred index maintenance.

Which three methods will remove the orphaned keys? (Choose three.)

A. using ALTER INDEX REBUILD PARTITION once for each affected local index partition

B. using DBMS\_PART.CLEANUP\_GIDX once to clean up orphaned keys in all local indexes.

C. using ALTER INDEX COALESCE PARTITION CLEANUP once for each affected global index partition

D. using DBMS\_PART.CLEANUP\_CIDX once to clean up any orphaned keys in all global indexes on partitioned tables.

E. using ALTER INDEX COALESCE PARTITION CLEANUP once for each affected local index partition

F. using ALTER INDEX REBUILD PARTITION once for each affected global index partition

Correct Answer: ABC