

# Oracle

## Exam 1z0-060

### Upgrade to Oracle Database 12c

Version: Demo

[ Total Questions: 10 ]

**Question No : 1**

Examine the following commands for redefining a table with Virtual Private Database (VPD) policies:

```
BEGIN
  DBMS_RLS.ADD_POLICY (
    object_schema => 'hr',
    object_name   => 'employees',
    policy_name   => 'employees_policy',
    function_schema => 'hr',
    policy_function => 'auth_emp_dep_100',
    statement_types => 'select, insert, update, delete'
  );
END;

BEGIN
  DBMS_REDEFINITION.START_REDEF_TABLE (
    uname          => 'hr',
    orig_table     => 'employees',
    int_table      => 'int_employees',
    col_mapping    => NULL,
    options_flag   => DBMS_REDEFINITION.CONSTRAINT_USE_PK,
    orderby_cols  => NULL,
    part_name      => NULL,
    copy_vpd_opt  => DBMS_REDEFINITION.CONSTRAINT_VPD_AUTO);
END;
```

Which two statements are true about redefining the table?

- A. All the triggers for the table are disabled without changing any of the column names or column types in the table.
- B. The primary key constraint on the EMPLOYEES table is disabled during redefinition.
- C. VPD policies are copied from the original table to the new table during online redefinition.
- D. You must copy the VPD policies manually from the original table to the new table during online redefinition.

**Answer: A,C**

**Explanation:** The triggers cloned to the interim table are disabled until the redefinition is completed. Once the redefinition is complete, all cloned objects are renamed to the original names used by the objects they were cloned from.

Ref: <http://www.oracle-base.com/articles/10g/online-table-redefinition-enhancements-10gr1.php>

**Question No : 2**

You notice a performance change in your production Oracle database and you want to know which change has made this performance difference.

You generate the Compare Period Automatic Database Diagnostic Monitor (ADDM) report to further investigation.

Which three findings would you get from the report?

- A. It detects any configuration change that caused a performance difference in both time periods.
- B. It identifies any workload change that caused a performance difference in both time periods.
- C. It detects the top wait events causing performance degradation.
- D. It shows the resource usage for CPU, memory, and I/O in both time periods.
- E. It shows the difference in the size of memory pools in both time periods.
- F. It gives information about statistics collection in both time periods.

**Answer: A,B,D**

**Explanation:** Keyword: shows the difference.

\* Full ADDM analysis across two AWR snapshot periods

Detects causes, measure effects, then correlates them

Causes: workload changes, configuration changes

Effects: regressed SQL, reach resource limits (CPU, I/O, memory, interconnect)

Makes actionable recommendations along with quantified impact

\* Identify what changed

/ Configuration changes, workload changes

\* Performance degradation of the database occurs when your database was performing optimally in the past, such as 6 months ago, but has gradually degraded to a point where it becomes noticeable to the users. The Automatic Workload Repository (AWR) Compare Periods report enables you to compare database performance between two periods of time.

While an AWR report shows AWR data between two snapshots (or two points in time), the AWR Compare Periods report shows the difference (ABE) between two periods (or two AWR reports with a total of four snapshots). Using the AWR Compare Periods report helps you to identify detailed performance attributes and configuration settings that differ between two time periods.

Reference: Resolving Performance Degradation Over Time

**Question No : 3**

Which two statements are true about the Oracle Direct Network File system (DNFS)?

- A. It utilizes the OS file system cache.
- B. A traditional NFS mount is not required when using Direct NFS.
- C. Oracle Disk Manager can manage NFS on its own, without using the operating kernel NFS driver.
- D. Direct NFS is available only in UNIX platforms.
- E. Direct NFS can load-balance I/O traffic across multiple network adapters.

**Answer: C,E**

**Explanation:** E: Performance is improved by load balancing across multiple network interfaces (if available).

Note:

\* To enable Direct NFS Client, you must replace the standard Oracle Disk Manager (ODM) library with one that supports Direct NFS Client.

Incorrect:

Not A: Direct NFS Client is capable of performing concurrent direct I/O, which bypasses any operating system level caches and eliminates any operating system write-ordering locks

Not B:

\* To use Direct NFS Client, the NFS file systems must first be mounted and available over regular NFS mounts.

\* Oracle Direct NFS (dNFS) is an optimized NFS (Network File System) client that provides faster and more scalable access to NFS storage located on NAS storage devices (accessible over TCP/IP).

Not D: Direct NFS is provided as part of the database kernel, and is thus available on all supported database platforms - even those that don't support NFS natively, like Windows.

Note:

\* Oracle Direct NFS (dNFS) is an optimized NFS (Network File System) client that provides faster and more scalable access to NFS storage located on NAS storage devices

(accessible over TCP/IP). Direct NFS is built directly into the database kernel - just like ASM which is mainly used when using DAS or SAN storage.

\* Oracle Direct NFS (dNFS) is an internal I/O layer that provides faster access to large NFS files than traditional NFS clients.

**Question No : 4**

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 for 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 controller failure group is brought back online, and newly allocated primary allocation units (AU) are stored in the controller2 failure group, without mirroring.
- E. Transactions accessing database objects contained in any tablespace stored in DGROUP1 will fail but queries will succeed.

**Answer: B**

**Explanation:** CREATE DISKGROUP NORMAL REDUNDANCY

\* For Oracle ASM to mirror files, specify the redundancy level as NORMAL REDUNDANCY (2-way mirroring by default for most file types) or HIGH REDUNDANCY (3-way mirroring for all files).

**Question No : 5**

You have installed two 64G flash devices to support the Database Smart Flash Cache feature on your database server that is running on Oracle Linux.

You have set the DB\_SMART\_FLASH\_FILE parameter:

```
DB_FLASH_CACHE_FILE= '/dev/flash_device_1 ',' /dev/flash_device_2'
```

How should the DB\_FLASH\_CACHE\_SIZE be configured to use both devices?

- A. Set DB\_FLASH\_CACHE\_SIZE = 64G.
- B. Set DB\_FLASH\_CACHE\_SIZE = 64G, 64G
- C. Set DB\_FLASH\_CACHE\_SIZE = 128G.
- D. DB\_FLASH\_CACHE\_SIZE is automatically configured by the instance at startup.

**Answer: B**

**Explanation:** \* Smart Flash Cache concept is not new in Oracle 12C - DB Smart Flash Cache in Oracle 11g.

In this release Oracle has made changes related to both initialization parameters used by DB Smart Flash cache. Now you can define many files|devices and its sizes for “Database Smart Flash Cache” area. In previous releases only one file|device could be defined.

```
DB_FLASH_CACHE_FILE = /dev/sda, /dev/sdb, /dev/sdc
```

```
DB_FLASH_CACHE_SIZE = 32G, 32G, 64G
```

So above settings defines 3 devices which will be in use by “DB Smart Flash Cache”

/dev/sda – size 32G

/dev/sdb – size 32G

/dev/sdc – size 64G

New view V\$FLASHFILESTAT – it’s used to determine the cumulative latency and read counts of each file|device and compute the average latency

**Question No : 6**

Your database is open and the LISTENER listener running. You stopped the wrong listener LISTENER by issuing the following command:

```
1snrctl > STOP
```

What happens to the sessions that are presently connected to the database Instance?

- A. They are able to perform only queries.
- B. They are not affected and continue to function normally.
- C. They are terminated and the active transactions are rolled back.
- D. They are not allowed to perform any operations until the listener LISTENER is started.

**Answer: B**

**Explanation:** The listener is used when the connection is established. The immediate impact of stopping the listener will be that no new session can be established from a remote host. Existing sessions are not compromised.

**Question No : 7**

Which three statements are true concerning the multitenant architecture?

- A. Each pluggable database (PDB) has its own set of background processes.
- B. A PDB can have a private temp tablespace.
- C. PDBs can share the sysaux tablespace.
- D. Log switches occur only at the multitenant container database (CDB) level.
- E. Different PDBs can have different default block sizes.
- F. PDBs share a common system tablespace.
- G. Instance recovery is always performed at the CDB level.

**Answer: B,D,G**

**Explanation: B:**

\* A PDB would have its SYSTEM, SYSAUX, TEMP tablespaces. It can also contain other

user created tablespaces in it.

\* There is one default temporary tablespace for the entire CDB. However, you can create additional temporary tablespaces in individual PDBs.

D:

\* There is a single redo log and a single control file for an entire CDB

\* A log switch is the point at which the database stops writing to one redo log file and begins writing to another. Normally, a log switch occurs when the current redo log file is completely filled and writing must continue to the next redo log file.

G: instance recovery

The automatic application of redo log records to uncommitted data blocks when an database instance is restarted after a failure.

Incorrect:

Not A:

\* There is one set of background processes shared by the root and all PDBs. –

\* High consolidation density. The many pluggable databases in a single container database share its memory and background processes, letting you operate many more pluggable databases on a particular platform than you can single databases that use the old architecture.

Not C: There is a separate SYSAUX tablespace for the root and for each PDB.

Not F: There is a separate SYSTEM tablespace for the root and for each PDB. -

**Question No : 8**

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.

**Answer: B**

**Explanation:** \* 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 No : 9**

After implementing full Oracle Data Redaction, you change the default value for the NUMBER data type as follows:

```
SQL> SELECT NUMBER_VALUE FROM REDACTION_VALUES_FOR_TYPE_FULL;
NUMBER_VALUE
-----
0

SQL> EXEC DBMS_REDACT.UPDATE_FULL_REDACTION_VALUES(-1)
PL/SQL procedure successfully completed.

SQL> select number_value from redaction_values_for_type_full;
NUMBER_VALUE
-----
-1
```

After changing the value, you notice that FULL redaction continues to redact numeric data with zero.

What must you do to activate the new default value for numeric full redaction?

- A. Re-enable redaction policies that use FULL data redaction.
- B. Re-create redaction policies that use FULL data redaction.
- C. Re-connect the sessions that access objects with redaction policies defined on them.
- D. Flush the shared pool.
- E. Restart the database instance.

**Answer: E**

**Explanation:** About Altering the Default Full Data Redaction Value

You can alter the default displayed values for full Data Redaction policies. By default, 0 is the redacted value when Oracle Database performs full redaction (DBMS\_REDACT.FULL) on a column of the NUMBER data type. If you want to change it to another value (for example, 7), then you can run the DBMS\_REDACT.UPDATE\_FULL\_REDACTION\_VALUES procedure to modify this value. The modification applies to all of the Data Redaction policies in the current database instance. After you modify a value, you must restart the database for it to take effect.

Note:

\* The DBMS\_REDACT package provides an interface to Oracle Data Redaction, which enables you to mask (redact) data that is returned from queries issued by low-privileged users or an application.

\* UPDATE\_FULL\_REDACTION\_VALUES Procedure

This procedure modifies the default displayed values for a Data Redaction policy for full redaction.

\* After you create the Data Redaction policy, it is automatically enabled and ready to redact data.

\* Oracle Data Redaction enables you to mask (redact) data that is returned from queries issued by low-privileged users or applications. You can redact column data by using one of the following methods:

/ Full redaction.

/ Partial redaction.

- / Regular expressions.
- / Random redaction.
- / No redaction.

Reference: Oracle Database Advanced Security Guide 12c, About Altering the Default Full Data Redaction Value

**Question No : 10**

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 outcome and actions to aggregate the 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. The trace files may be aggregated by using the trcess utility.
5. The trace files be aggregated by using the tkprof utility.

Identify the correct outcome and the step to aggregate by using tkprof utility?

- A. 1
- B. 2 and 4
- C. 2 and 5
- D. 3 and 4
- E. 3 and 5

**Answer: B**

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

Incorrect:

Not 1: Parameter `service_name`

Name of the service for which tracing is enabled.

`module_name`

Name of the MODULE. An optional additional qualifier for the service.

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);
```