



Oracle

1Z0-060 Exam

Upgrade to Oracle Database 12c Exam

Thank you for Downloading 1Z0-060 exam PDF Demo

You can also try our 1Z0-060 practice exam software

Download Free Demo:

<https://www.dumpsschool.com/1Z0-060-exam-dumps.html>

**Questions & Answers
(Demo Version – Limited Content)**

Question: 1

Your multitenant container (CDB) contains two pluggable databases (PDB), HR_PDB and ACCOUNTS_PDB, both of which use the CDB tablespace. The temp file is called temp01.tmp.

A user issues a query on a table on one of the PDBs and receives the following error:

ERROR at line 1:

ORA-01565: error in identifying file '/u01/app/oracle/oradata/CDB1/temp01.tmp'

ORA-27037: unable to obtain file status

Identify two ways to rectify the error.

- A. Add a new temp file to the temporary tablespace and drop the temp file that that produced the error.
- B. Shut down the database instance, restore the temp01.tmp file from the backup, and then restart the database.
- C. Take the temporary tablespace offline, recover the missing temp file by applying redo logs, and then bring the temporary tablespace online.
- D. Shutdown the database instance, restore and recover the temp file from the backup, and then open the database with RESETLOGS.
- E. Shut down the database instance and then restart the CDB and PDBs.

Answer: CE

Explanation:

* Because temp files cannot be backed up and because no redo is ever generated for them, RMAN never restores or recovers temp files. RMAN does track the names of temp files, but only so that it can automatically re-create them when needed.

* If you use RMAN in a Data Guard environment, then RMAN transparently converts primary control files to standby control files and vice versa. RMAN automatically updates file names for data files, online redo logs, standby redo logs, and temp files when you issue RESTORE and RECOVER.

Question: 2

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.CONST_USE_PK,
    orderby_cols  => NULL,
    part_name      => NULL,
    copy_vpd_opt  => DBMS_REDEFINITION.CONST_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: BC

Explanation:

C (not D): CONS_VPD_AUTO

Used to indicate to copy VPD policies automatically

* DBMS_RLS.ADD_POLICY

/ The DBMS_RLS package contains the fine-grained access control administrative interface, which is used to implement Virtual Private Database (VPD). DBMS_RLS is available with the Enterprise Edition only.

Note:

* CONS_USE_PK and CONS_USE_ROWID are constants used as input to the "options_flag" parameter in both the START_REDEF_TABLE Procedure and CAN_REDEF_TABLE Procedure. CONS_USE_ROWID is used to indicate that the redefinition should be done using rowids while CONS_USE_PK implies that the redefinition should be done using primary keys or pseudo-primary keys (which are unique keys with all component columns having NOT NULL constraints).

* DBMS_REDEFINITION.START_REDEF_TABLE

To achieve online redefinition, incrementally maintainable local materialized views are used. These logs keep track of the changes to the master tables and are used by the materialized views during refresh synchronization.

* START_REDEF_TABLE Procedure

Prior to calling this procedure, you must manually create an empty interim table (in the same schema as the table to be redefined) with the desired attributes of the post-redefinition table, and then call this procedure to initiate the redefinition.

Question: 3

Which two statements are true about the use of the procedures listed in the `v$sysaux_occupants.move_procedure` column?

- A. The procedure may be used for some components to relocate component data to the SYSAUX tablespace from its current tablespace.
- B. The procedure may be used for some components to relocate component data from the SYSAUX tablespace to another tablespace.
- C. All the components may be moved into SYSAUX tablespace.
- D. All the components may be moved from the SYSAUX tablespace.

Answer: BD

Explanation:

`V$SYSAUX_OCCUPANTS` displays SYSAUX tablespace occupant information.

`MOVE_PROCEDURE`: Name of the move procedure; null if not applicable

For example, the tables and indexes that were previously owned by the system user can now be specified for a SYSAUX tablespace. You can query the `v$sysaux_occupants` view to find the exact components stored within the SYSAUX tablespace.

Question: 4

Which statement is true about Oracle Net Listener?

- A. It acts as the listening endpoint for the Oracle database instance for all local and non-local user connections.
- B. A single listener can service only one database instance and multiple remote client connections.
- C. Service registration with the listener is performed by the process monitor (PMON) process of each database instance.
- D. The listener.ora configuration file must be configured with one or more listening protocol addresses to allow remote users to connect to a database instance.
- E. The listener.ora configuration file must be located in the `ORACLE_HOME/network/admin` directly.

Answer: C

Explanation:

Supported services, that is, the services to which the listener forwards client requests, can be configured in the listener.ora file or this information can be dynamically registered with the listener. This dynamic registration feature is called service registration. The registration is performed by the PMON process—an instance background process—of each database instance that has the necessary configuration in the database initialization parameter file. Dynamic service registration does not require any configuration in the listener.ora file.

Incorrect:

Not B: Service registration reduces the need for the SID_LIST_listener_name parameter setting, which specifies information about the databases served by the listener, in the listener.ora file.

Note:

* Oracle Net Listener is a separate process that runs on the database server computer. It receives incoming client connection requests and manages the traffic of these requests to the database server.

* A remote listener is a listener residing on one computer that redirects connections to a database instance on another computer. Remote listeners are typically used in an Oracle Real Application Clusters (Oracle RAC) environment. You can configure registration to remote listeners, such as in the case of Oracle RAC, for dedicated server or shared server environments.

Question: 5

You are administering a database stored in Automatic Storage Management (ASM). You use RMAN to back up the database and the MD_BACKUP command to back up the ASM metadata regularly. You lost an ASM disk group DG1 due to hardware failure. In which three ways can you re-create the lost disk group and restore the data?

- A. Use the MD_RESTORE command to restore metadata for an existing disk group by passing the existing disk group name as an input parameter and use RMAN to restore the data.
- B. Use the MKDGG command to restore the disk group with the same configuration as the backed-up disk group and data on the disk group.
- C. Use the MD_RESTORE command to restore the disk group with the changed disk group specification, failure group specification, name, and other attributes and use RMAN to restore the data.
- D. Use the MKDGG command to restore the disk group with the same configuration as the backed-up disk group name and same set of disks and failure group configuration, and use RMAN to restore the data.
- E. Use the MD_RESTORE command to restore both the metadata and data for the failed disk group.
- F. Use the MKDGG command to add a new disk group DG1 with the same or different specifications for failure group and other attributes and use RMAN to restore the data.

Answer: CEF

Explanation:

Note:

* The [md_restore](#) command allows you to restore a disk group from the metadata created by the md_backup command.

/md_restore Command

Purpose

This command restores a disk group backup using various options that are described in this section.

/ In the restore mode md_restore, it re-creates the disk group based on the backup file with all user-defined templates with the exact configuration as the backed-up disk group. There are several options when restoring the disk group

full - re-creates the disk group with the exact configuration

nodg - Restores metadata in an existing disk group provided as an input parameter

newdgg - Change the configuration like failure group, disk group name, etc..

* The MD_BACKUP command creates a backup file containing metadata for one or more disk groups. By default all the mounted disk groups are included in the backup file which is saved in the current working directory. If the name of the backup file is not specified, ASM names the file AMBR_BACKUP_INTERMEDIATE_FILE.

Question: 6

Your multitenant container database, CDB1, is running in ARCHIVELOG mode and has two pluggable databases, HR_PDB and ACCOUNTS_PDB. An RMAN backup exists for the database. You issue the command to open ACCOUNTS_PDB and find that the USERDATA.DBF data file for the default permanent tablespace USERDATA belonging to ACCOUNTS_PDB is corrupted.

What should you do before executing the commands to restore and recover the data file in ACCOUNTS_PDB?

- A. Place CDB1 in the mount stage and then the USERDATA tablespace offline in ACCOUNTS_PDB.
- B. Place CDB1 in the mount stage and issue the ALTER PLUGGABLE DATABASE accounts_pdb CLOSE IMMEDIATE command.
- C. Issue the ALTER PLUGGABLE DATABASE accounts_pdb RESTRICTED command.
- D. Take the USERDATA tablespace offline in ACCOUNTS_PDB.

Answer: D

Explanation:

* You can take an online tablespace offline so that it is temporarily unavailable for general use. The rest of the database remains open and available for users to access data. Conversely, you can bring an offline tablespace online to make the schema objects within the tablespace available to database users. The database must be open to alter the availability of a tablespace.

Question: 7

Which Oracle Database component is audited by default if the unified Auditing option is enabled?

- A. Oracle Data Pump
- B. Oracle Recovery Manager (RMAN)
- C. Oracle Label Security
- D. Oracle Database Vault
- E. Oracle Real Application Security

Answer: E

Explanation:

Type of Unified auditing:

Standard

Fine Grained Audit

XS

Database Vault (not D)

Label Security (not C)

RMAN AUDIT (not B)

Data Pump (not A)

Note:

* Oracle 12c introduces Unified Auditing, which consolidates database audit records including :-

DDL, DML, DCL

Fine Grained Auditing (DBMS_FGA)

Oracle Database Real Application Security

Oracle Recovery Manager

Oracle Database Vault

Oracle Label Security

Oracle Data Mining

Oracle Data Pump

Oracle SQL*Loader Direct Load

Question: 8

Your multitenant container (CDB) containing three pluggable databases (PDBs) is running in ARCHIVELOG mode. You find that the SYSAUX tablespace is corrupted in the root container. The steps to recover the tablespace are as follows:

1. Mount the CDB.
2. Close all the PDBs.
3. Open the database.
4. Apply the archive redo logs.
5. Restore the data file.
6. Take the SYSAUX tablespace offline.
7. Place the SYSAUX tablespace offline.
8. Open all the PDBs with RESETLOGS.
9. Open the database with RESETLOGS.
10. Execute the command SHUTDOWN ABORT.

Which option identifies the correct sequence to recover the SYSAUX tablespace?

- A. 6, 5, 4, 7
- B. 10, 1, 2, 5, 8
- C. 10, 1, 2, 5, 4, 9, 8
- D. 10, 1, 5, 8, 10

Answer: C

Explanation:

* Example:

While evaluating the 12c beta3 I was not able to do the recover while testing "all pdb files lost".

Cannot close the pdb as the system datafile was missing...

So only option to recover was:

Shutdown cdb (10)

startup mount; (1)

restore pluggable database

recover pluggable databsae

alter database open;

alter pluggable database name open;

Oracle support says: You should be able to close the pdb and restore/recover the system tablespace of PDB.

* Open the database with the RESETLOGS option after finishing recovery:

```
SQL> ALTER DATABASE OPEN RESETLOGS;
```

Question: 9

Which three are direct benefits of the multiprocess, multithreaded architecture of Oracle Database 12c when it is enabled?

- A. Reduced logical I/O
- B. Reduced virtual memory utilization
- C. Improved parallel Execution performance
- D. Improved Serial Execution performance
- E. Reduced physical I/O
- F. Reduced CPU utilization

Answer: C

Explanation:

* Multiprocess and Multithreaded Oracle Database Systems

Multiprocess Oracle Database (also called multiuser Oracle Database) uses several processes to run different parts of the Oracle Database code and additional Oracle processes for the users—either one process for each connected user or one or more processes shared by multiple users. Most databases are multiuser because a primary advantage of a database is managing data needed by multiple users simultaneously.

Each process in a database instance performs a specific job. By dividing the work of the database and applications into several processes, multiple users and applications can connect to an instance simultaneously while the system gives good performance.

* In previous releases, Oracle processes did not run as threads on UNIX and Linux systems. Starting in Oracle Database 12c, the multithreaded Oracle Database model enables Oracle processes to execute as operating system threads in separate address spaces.

Question: 10

In order to exploit some new storage tiers that have been provisioned by a storage administrator, the partitions of a large heap table must be moved to other tablespaces in your Oracle 12c database?

Both local and global partitioned B-tree Indexes are defined on the table.

A high volume of transactions access the table during the day and a medium volume of transactions access it at night and during weekends.

Minimal disruption to availability is required.

Which three statements are true about this requirement?

- A. The partitions can be moved online to new tablespaces.
- B. Global indexes must be rebuilt manually after moving the partitions.
- C. The partitions can be compressed in the same tablespaces.

- D. The partitions can be compressed in the new tablespaces.
- E. Local indexes must be rebuilt manually after moving the partitions.

Answer: ABD

Explanation:

A: You can create and rebuild indexes online. Therefore, you can update base tables at the same time you are building or rebuilding indexes on that table. You can perform DML operations while the index build is taking place, but DDL operations are not allowed. Parallel execution is not supported when creating or rebuilding an index online.

B:

Note:

* Transporting and Attaching Partitions for Data Warehousing Typical enterprise data warehouses contain one or more large fact tables. These fact tables can be partitioned by date, making the enterprise data warehouse a historical database. You can build indexes to speed up star queries. Oracle recommends that you build local indexes for such historically partitioned tables to avoid rebuilding global indexes every time you drop the oldest partition from the historical database.

D: Moving (Rebuilding) Index-Organized Tables

Because index-organized tables are primarily stored in a B-tree index, you can encounter fragmentation as a consequence of incremental updates. However, you can use the ALTER TABLE...MOVE statement to rebuild the index and reduce this fragmentation.

Question: 11

Which three are true about the large pool for an Oracle database instance that supports shared server connections?

- A. Allocates memory for RMAN backup and restore operations
- B. Allocates memory for shared and private SQL areas
- C. Contains a cursor area for storing runtime information about cursors
- D. Contains stack space
- E. Contains a hash area performing hash joins of tables

Answer: ABC

Explanation:

The large pool can provide large memory allocations for the following:

/ (B)UGA (User Global Area) for the shared server and the Oracle XA interface (used where transactions interact with multiple databases)

/Message buffers used in the parallel execution of statements

/ (A) Buffers for Recovery Manager (RMAN) I/O slaves

Note:

* large pool

Optional area in the SGA that provides large memory allocations for backup and restore operations,

I/O server processes, and session memory for the shared server and Oracle XA.

* Oracle XA

An external interface that allows global transactions to be coordinated by a transaction manager other than Oracle Database.

* UGA

User global area. Session memory that stores session variables, such as logon information, and can also contain the OLAP pool.

* Configuring the Large Pool

Unlike the shared pool, the large pool does not have an LRU list (not D). Oracle Database does not attempt to age objects out of the large pool. Consider configuring a large pool if the database instance uses any of the following Oracle Database features:

* Shared server

In a shared server architecture, the session memory for each client process is included in the shared pool.

* Parallel query

Parallel query uses shared pool memory to cache parallel execution message buffers.

* Recovery Manager

Recovery Manager (RMAN) uses the shared pool to cache I/O buffers during backup and restore operations. For I/O server processes, backup, and restore operations, Oracle Database allocates buffers that are a few hundred kilobytes in size.

Question: 12

What are three purposes of the RMAN "FROM" clause?

- A. to support PUSH-based active database duplication
- B. to support synchronization of a standby database with the primary database in a Data environment
- C. To support PULL-based active database duplication
- D. To support file restores over the network in a Data Guard environment
- E. To support file recovery over the network in a Data Guard environment

Answer: BCE

Explanation:

E:

* With a control file autobackup, RMAN can recover the database even if the current control file, recovery catalog, and server parameter file are inaccessible.

* RMAN uses a recovery catalog to track filenames for all database files in a Data Guard environment. A recovery catalog is a database schema used by RMAN to store metadata about one or more Oracle databases. The catalog also records where the online redo logs, standby redo logs, tempfiles, archived redo logs, backup sets, and image copies are created.

Question: 13

You notice that the performance of your production 24/7 Oracle database significantly degraded. Sometimes you are not able to connect to the instance because it hangs. You do not want to restart

the database instance. How can you detect the cause of the degraded performance?

- A. Enable Memory Access Mode, which reads performance data from SGA.
- B. Use emergency monitoring to fetch data directly from SGA analysis.
- C. Run Automatic Database Diagnostic Monitor (ADDM) to fetch information from the latest Automatic Workload Repository (AWR) snapshots.
- D. Use Active Session History (ASH) data and hang analysis in regular performance monitoring.
- E. Run ADDM in diagnostic mode.

Answer: C

Explanation:

* In most cases, ADDM output should be the first place that a DBA looks when notified of a performance problem.

* 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 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: 14

You plan to use the In Database Archiving feature of Oracle Database 12c, and store rows that are inactive for over three months, in Hybrid Columnar Compressed (HCC) format. Which three storage options support the use of HCC?

- A. ASM disk groups with ASM disks consisting of Exadata Grid Disks.
- B. ASM disk groups with ASM disks consisting of LUNs on any Storage Area Network array
- C. ASM disk groups with ASM disks consisting of any zero padded NFS-mounted files
- D. Database files stored in ZFS and accessed using conventional NFS mounts.
- E. Database files stored in ZFS and accessed using the Oracle Direct NFS feature
- F. Database files stored in any file system and accessed using the Oracle Direct NFS feature
- G. ASM disk groups with ASM disks consisting of LUNs on Pillar Axiom Storage arrays

Answer: AEG

Explanation:

HCC requires the use of Oracle Storage – Exadata (A), Pillar Axiom (G) or Sun ZFS Storage Appliance (ZFSSA).

Note:

* Hybrid Columnar Compression, initially only available on Exadata, has been extended to support Pillar Axiom and Sun ZFS Storage Appliance (ZFSSA) storage when used with Oracle Database Enterprise Edition 11.2.0.3 and above

* Oracle offers the ability to manage NFS using a feature called Oracle Direct NFS (dNFS). Oracle Direct NFS implements NFS V3 protocol within the Oracle database kernel itself. Oracle Direct NFS client overcomes many of the challenges associated with using NFS with the Oracle Database with simple configuration, better performance than traditional NFS clients, and offers consistent configuration across platforms.

Question: 15

In your multitenant container database (CDB) containing pluggable databases (PDB), users complain about performance degradation.

How does real-time Automatic database Diagnostic Monitor (ADDM) check performance degradation and provide solutions?

- A. It collects data from SGA and compares it with a preserved snapshot.
- B. It collects data from SGA, analyzes it, and provides a report.
- C. It collects data from SGA and compares it with the latest snapshot.
- D. It collects data from both SGA and PGA, analyzes it, and provides a report.

Answer: B

Explanation:

Note:

* The multitenant architecture enables an Oracle database to function as a multitenant container database (CDB) that includes zero, one, or many customer-created pluggable databases (PDBs). A PDB is a portable collection of schemas, schema objects, and nonschema objects that appears to an Oracle Net client as a non-CDB. All Oracle databases before Oracle Database 12c were non-CDBs.

* The System Global Area (SGA) is a group of shared memory areas that are dedicated to an Oracle "instance" (an instance is your database programs and RAM).

* The PGA (Program or Process Global Area) is a memory area (RAM) that stores data and control information for a single process.

Thank You for trying 1Z0-060 PDF Demo

To try our 1Z0-060 practice exam software visit link below

<https://www.dumpsschool.com/1Z0-060-exam-dumps.html>

Start Your 1Z0-060 Preparation

Use Coupon “20OFF” for extra 20% discount on the purchase of Practice Test Software. Test your 1Z0-060 preparation with actual exam questions.