



BrainDumps  
Collection

# Oracle

## 1Z0-060 Exam

Upgrade to Oracle Database 12c

Thank you for Downloading 1Z0-060 exam PDF Demo

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

**Download Free Demo**

<https://www.braindumpscollection.com/1Z0-060.html>

DEMO  
VERSION (LIMITED CONTENT)

Questions  
& Answers

# Version: 14.0

---

**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. (Choose two.)

- 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. Shut down 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: A,E**

---

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.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? (Choose two.)

- 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 object they were cloned from.

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

---

### Question: 3

---

Which two statements are true about the use of the procedures listed in the v\$sysaux\_occupants.move\_procedure column? (Choose two.)

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

---

**Answer: A,B**

---

Explanation:

References: [http://www.dba-oracle.com/t\\_v\\_sysaux\\_contents\\_tips.htm](http://www.dba-oracle.com/t_v_sysaux_contents_tips.htm)

---

#### **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 listener registration process (LREG) 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 directory.

---

**Answer: C**

---

Explanation:

---

#### **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? (Choose three.)

- 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 MKDG 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 MKDG 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: C,D,F**

---

Explanation:

---

**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 take 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 OPEN 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: B**

---

Explanation:

---

**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 online.
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: A**

---

Explanation:

```
RMAN> ALTER TABLESPACE sysaux OFFLINE IMMEDIATE;
```

```
RMAN> RESTORE TABLESPACE sysaux;
```

```
RMAN> RECOVER TABLESPACE sysaux;
```

```
RMAN> ALTER TABLESPACE sysaux ONLINE;
```

\* 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 database
```

```
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.

\* Inconsistent backups are usually created by taking online database backups. You can also make an inconsistent backup by backing up data files while a database is closed, either:

/ Immediately after the crash of an Oracle instance (or, in an Oracle RAC configuration, all instances)

/ After shutting down the database using SHUTDOWN ABORT

Inconsistent backups are only useful if the database is in ARCHIVELOG mode and all archived redo logs created since the backup are available.

\* 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? (Choose three.)

- 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: B,C,F**

---

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? (Choose three.)

- 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: A,C,D**

---

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.

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.

C: If a table can be compressed in the new tablespace, also it can be compressed in the same tablespace.

---

### Question: 11

---

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

- 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: A,B,C**

---

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? (Choose three.)

- 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: B,D,E**

---

Explanation:

- \* 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 12c database has 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 for 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: B**

---

Explanation:

---

**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? (Choose three.)

- 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: A,E,G**

---

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.

---

**Question: 16**

---

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.

---

**Answer: C**

---

Explanation:

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: 17**

---

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

---

**Answer: B**

---

Explanation:

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: 18

---

Your database is running an ARCHIVELOG mode.

The following parameters are set in your database instance:

```
LOG_ARCHIVE_FORMAT = arch %t %s %r.arc  
LOG_ARCHIVE_DEST_1 = 'LOCATION = /disk1/archive'  
DB_RECOVERY_FILE_DEST_SIZE = 50G  
DB_RECOVERY_FILE = '/u01/oradata'
```

Which statement is true about the archived redo log files?

- A. They are created only in the location specified by the LOG\_ARCHIVE\_DEST\_1 parameter.
- B. They are created only in the Fast Recovery Area because configuring the DB\_RECOVERY\_FILE\_DEST and DB\_RECOVERY\_FILE\_DEST\_SIZE parameters automatically enables flashback for the database.
- C. They are created in the location specified by the LOG\_ARCHIVE\_DEST\_1 parameter and in the default location \$ORACLE\_HOME/dbs/arch.
- D. They are created in the location specified by the LOG\_ARCHIVE\_DEST\_1 parameter and in the location specified by the DB\_RECOVERY\_FILE\_DEST parameter.

---

**Answer: A**

---

Explanation:

You can choose to archive redo logs to a single destination or to multiple destinations.

Destinations can be local—within the local file system or an Oracle Automatic Storage Management (Oracle ASM) disk group—or remote (on a standby database). When you archive to multiple

destinations, a copy of each filled redo log file is written to each destination. These redundant copies help ensure that archived logs are always available in the event of a failure at one of the destinations. To archive to only a single destination, specify that destination using the LOG\_ARCHIVE\_DEST and LOG\_ARCHIVE\_DUPLEX\_DEST initialization parameters.

ARCHIVE\_DEST initialization parameter. To archive to multiple destinations, you can choose to archive to two or more locations using the LOG\_ARCHIVE\_DEST\_n initialization parameters, or to archive only to a primary and secondary destination using the LOG\_ARCHIVE\_DEST and LOG\_ARCHIVE\_DUPLEX\_DEST initialization parameters.

---

**Question: 19**

---

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

---

**Answer: B**

---

Explanation:

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: 20**

---

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.

---

**Answer: C**

---

Explanation:

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.

## Thank You for trying 1Z0-060 PDF Demo

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

<https://www.braindumpscollection.com/1Z0-060.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.