


1Z0-062 - Oracle Database 12c: Installation and Administration

<https://www.certleader.com/1Z0-062-dumps.html>



1. You are required to migrate your 11.2.0.3 database as a pluggable database (PDB) to a multitenant container database (CDB).

The following are the possible steps to accomplish this task:

1. Place all the user-defined tablespaces in read-only mode on the source database.
2. Upgrade the source database to a 12c version.
3. Create a new PDB in the target container database.
4. Perform a full transportable export on the source database with the VERSION parameter set to 12 using the expdp utility.
5. Copy the associated data files and export the dump file to the desired location in the target database.
6. Invoke the Data Pump import utility on the new PDB database as a user with the DATAPUMP_IMP_FULL_DATABASE role and specify the full transportable import options.
7. Synchronize the PDB on the target container database by using the DBMS_PDS.SYNC_ODB function.

Identify the correct order of the required steps.

- A. 2, 1, 3, 4, 5, 6
- B. 1, 3, 4, 5, 6, 7
- C. 1, 4, 3, 5, 6, 7
- D. 2, 1, 3, 4, 5, 6, 7
- E. 1, 5, 6, 4, 3, 2

Answer: C

Explanation:

1. Set user tablespaces in the source database to READ ONLY.
2. From the Oracle Database 11g Release 2 (11.2.0.3) environment, export the metadata and any data residing in administrative tablespaces from the source database using the FULL=Y and TRANSPORTABLE=ALWAYS parameters. Note that the VERSION=12 parameter is required only when exporting from an Oracle Database 11g Release 2 database:
3. Copy the tablespace data files from the source system to the destination system. Note that the log file from the export operation will list the data files required to be moved.
4. Create a CDB on the destination system, including a PDB into which you will import the source database.
5. In the Oracle Database 12c environment, connect to the pre-created PDB and import the dump file. The act of importing the dump file will plug the tablespace data files into the destination PDB Oracle White Paper - Upgrading to Oracle Database 12c -August 2013

2. You enabled an audit policy by issuing the following statements:

```
SQL> AUDIT POLICY ORA_DATABASE_PARAMETER BY SCOTT;
```

```
SQL> AUDIT POLICY ORA_DATABASE_PARAMETER BY SYS, SYSTEM;
```

For which database users and for which executions is the audit policy now active? Select two.

- A. SYS, SYSTEM
- B. SCOTT
- C. Only for successful executions
- D. Only for failed executions
- E. Both successful and failed executions

Answer: A,E

Explanation: * The ORA_DATABASE_PARAMETER policy audits commonly used Oracle Database parameter settings. By default, this policy is not enabled.

3. You must track all transactions that modify certain tables in the sales schema for at least three years.

Automatic undo management is enabled for the database with a retention of one day.

Which two must you do to track the transactions?

- A. Enable supplemental logging for the database.
- B. Specify undo retention guarantee for the database.
- C. Create a Flashback Data Archive in the tablespace where the tables are stored.
- D. Create a Flashback Data Archive in any suitable tablespace.
- E. Enable Flashback Data Archiving for the tables that require tracking.

Answer: D,E

Explanation: E: By default, flashback archiving is disabled for any table. You can enable flashback archiving for a table if you have the FLASHBACK ARCHIVE object privilege on the Flashback Data Archive that you want to use for that table.

D: Creating a Flashback Data Archive

/ Create a Flashback Data Archive with the CREATE FLASHBACK ARCHIVE statement, specifying the following:

Name of the Flashback Data Archive

Name of the first tablespace of the Flashback Data Archive

(Optional) Maximum amount of space that the Flashback Data Archive can use in the first tablespace

/ Create a Flashback Data Archive named fla2 that uses tablespace tbs2, whose data will be retained for two years:

```
CREATE FLASHBACK ARCHIVE fla2 TABLESPACE tbs2 RETENTION 2 YEAR;
```

4. A warehouse fact table in your Oracle 12c Database is range-partitioned by month and accessed frequently with queries that span multiple partitions

The table has a local prefixed, range partitioned index.

Some of these queries access very few rows in some partitions and all the rows in other partitions, but these queries still perform a full scan for all accessed partitions.

This commonly occurs when the range of dates begins at the end of a month or ends close to the start of a month.

You want an execution plan to be generated that uses indexed access when only a few rows are accessed from a segment, while still allowing full scans for segments where many rows are returned.

Which three methods could transparently help to achieve this result?

- A. Using a partial local Index on the warehouse fact table month column with indexing disabled to the table partitions that return most of their rows to the queries.
- B. Using a partial local Index on the warehouse fact table month column with indexing disabled for the table partitions that return a few rows to the queries.
- C. Using a partitioned view that does a UNION ALL query on the partitions of the warehouse fact table, which retains the existing local partitioned column.
- D. Converting the partitioned table to a partitioned view that does a UNION ALL query on the monthly tables, which retains the existing local partitioned column.
- E. Using a partial global index on the warehouse fact table month column with indexing disabling for the table partitions that return most of their rows to the queries.
- F. Using a partial global index on the warehouse fact table month column with indexing disabled for the table partitions that return a few rows to the queries.

Answer: B,C,E

Explanation:

Note:

* Oracle 12c now provides the ability to index a subset of partitions and to exclude the others.

Local and global indexes can now be created on a subset of the partitions of a table. Partial Global indexes provide more flexibility in index creation for partitioned tables. For example, index segments can be omitted for the most recent partitions to ensure maximum data ingest rates without impacting the overall data model and access for the partitioned object.

Partial Global Indexes save space and improve performance during loads and queries. This feature supports global indexes that include or index a certain subset of table partitions or subpartitions, and exclude the others. This operation is supported using a default table indexing property. When a table is created or altered, a default indexing property can be specified for the table or its partitions.

5. Your multitenant container database has three pluggable databases (PDBs): PDB1, PDB2, and PDB3.

Which two RMAN commands may be; used to back up only the PDB1 pluggable database?

- A. BACKUP PLUGGABLE DATABASE PDB1 while connected to the root container
- B. BACKUP PLUGGABLE DATABASE PDB1 while connected to the PDB1 container
- C. BACKUP DATABASE while connected to the PDB1 container
- D. BACKUP DATABASE while connected to the boot container
- E. BACKUP PLUGGABLE database PDB1 while connected to PDB2

Answer: A,C

Explanation: To perform operations on a single PDB, you can connect as target either to the root or directly to the PDB.

* (A) If you connect to the root, you must use the PLUGGABLE DATABASE syntax in your RMAN commands. For example, to back up a PDB, you use the BACKUP PLUGGABLE DATABASE command.

* (C) If instead you connect directly to a PDB, you can use the same commands that you would use when connecting to a non-CDB. For example, to back up a PDB, you would use the BACKUP DATABASE command.

Reference: Oracle Database Backup and Recovery User's Guide 12c, About Backup and Recovery of CDBs

6. You executed a DROP USER CASCADE on an Oracle 11g release 1 database and immediately realized that you forgot to copy the OCA.EXAM_RESULTS table to the OCP schema.

The RECYCLE_BIN enabled before the DROP USER was executed and the OCP user has been granted the FLASHBACK ANY TABLE system privilege.

What is the quickest way to recover the contents of the OCA.EXAM_RESULTS table to the OCP schema?

- A. Execute FLASHBACK TABLE OCA.EXAM_RESULTS TO BEFORE DROP RENAME TO OCP.EXAM_RESULTS; connected as SYSTEM.

- B. Recover the table using traditional Tablespace Point In Time Recovery.
- C. Recover the table using Automated Tablespace Point In Time Recovery.
- D. Recovery the table using Database Point In Time Recovery.
- E. Execute FLASHBACK TABLE OCA.EXAM_RESULTS TO BEFORE DROP RENAME TO EXAM_RESULTS; connected as the OCP user.

Answer: C

Explanation: RMAN tablespace point-in-time recovery (TSPITR).

Recovery Manager (RMAN) TSPITR enables quick recovery of one or more tablespaces in a database to an earlier time without affecting the rest of the tablespaces and objects in the database.

Fully Automated (the default)

In this mode, RMAN manages the entire TSPITR process including the auxiliary instance.

You specify the tablespaces of the recovery set, an auxiliary destination, the target time, and you allow RMAN to manage all other aspects of TSPITR.

The default mode is recommended unless you specifically need more control over the location of recovery set files after TSPITR, auxiliary set files during TSPITR, channel settings and parameters or some other aspect of your auxiliary instance.

7. 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

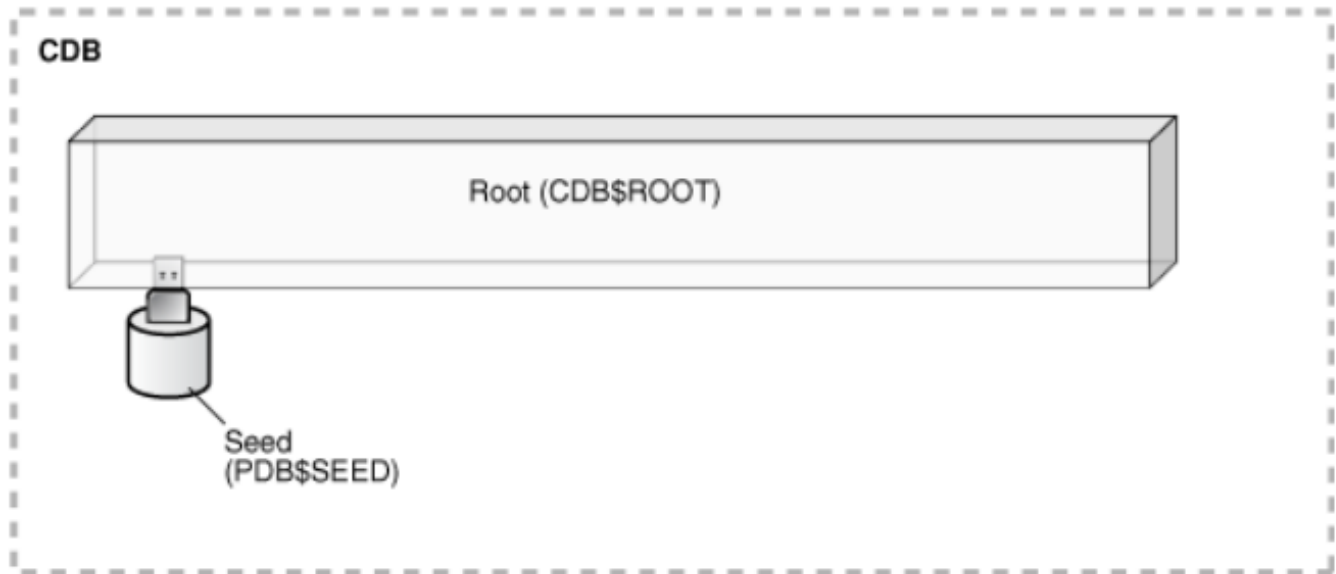
8. What is the effect of specifying the "ENABLE PLUGGABLE DATABASE" clause in a "CREATE DATABASE" statement?

- 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.

Answer: B

Explanation: * 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:



Description of admin095.png follows

* 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.

9. Examine the structure of the sales table, which is stored in a locally managed tablespace with Automatic Segment Space Management (ASSM) enabled.

NameNull?Type

PROD_IDNOT NULL NUMBER

CUST_IDNOT NULL NUMBER

TIME_IDNOT NULL DATE

CHANNEL_IDNOT NULL NUMBER

PROMO_IDNOT NULL NUMBER

QUANTITY__SOLDNOT NULL NUMBER (10, 2) AMOUNT_SOLDNOT NULL NUMBER (10, 2)

You want to perform online segment shrink to reclaim fragmented free space below the high water mark.

What should you ensure before the start of the operation?

- A. Row movement is enabled.
- B. Referential integrity constraints for the table are disabled.

- C. No queries are running on this table.
- D. Extra disk space equivalent to the size of the segment is available in the tablespace.
- E. No pending transaction exists on the table.

Answer: D

10. Which two statements are true about the logical storage structure of an Oracle database?

- A. An extent contains data blocks that are always physically contiguous on disk.
- B. An extent can span multiple segments,
- C. Each data block always corresponds to one operating system block.
- D. It is possible to have tablespaces of different block sizes.
- E. A data block is the smallest unit of I/O in data files.

Answer: B,D

Reference: http://docs.oracle.com/cd/E11882_01/server.112/e40540/logical.htm#CNCPT250

11. Your multitenant container database (CDB) contains pluggable databases (PDBs), you are connected to the HR_PDB. You execute the following command:

```
SQL > CREATE UNDO TABLESPACE undotb01
```

```
DATAFILE 'u01/oracle/rddb1/undotbs01.dbf' SIZE 60M AUTOEXTEND ON;
```

What is the result?

- A. It executes successfully and creates an UNDO tablespace in HR_PDB.
- B. It fails and reports an error because there can be only one undo tablespace in a CDB.
- C. It fails and reports an error because the CONTAINER=ALL clause is not specified in the command.
- D. It fails and reports an error because the CONTAINER=CURRENT clause is not specified in the command.
- E. It executes successfully but neither tablespace nor the data file is created.

Answer: E

Explanation: Interesting behavior in 12.1.0.1 DB of creating an undo tablespace in a PDB. With the new Multitenant architecture the undo tablespace resides at the CDB level and PDBs all share the same UNDO tablespace.

When the current container is a PDB, an attempt to create an undo tablespace fails without returning an

error.

12. Your multitenant container database (CDB) contains some pluggable databases (PDBs), you execute the following command in the root container:

```
SQL> CREATE USER c##a_admin  
IDENTIFIED BY password  
DEFAULT TABLESPACE data_ts  
QUOTA 100M ON test_ts  
QUOTA 500K ON data_ts  
TEMPORARY TABLESPACE temp_ts  
PROFILE hr_profile;
```

Which two statements are true?

- A. Schema objects owned by the C##A_ADMIN common user can be shared across all PDBs.
- B. The C##A_ADMIN user will be able to use the TEMP_TS temporary tablespace only in root.
- C. The command will, create a common user whose description is contained in the root and each PDB.
- D. The schema for the common user C##A_ADMIN can be different in each container.
- E. The command will create a user in the root container only because the container clause is not used.

Answer: C,D

13. 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. -

14. Which three factors influence the optimizer's choice of an execution plan?

A. the optimizer_mode initialization parameter

B. operating system (OS) statistics

C. cardinality estimates

D. object statistics in the data dictionary

E. fixed baselines

Answer: A,B

Reference: http://docs.oracle.com/cd/B19306_01/server.102/b14211/optimops.htm

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 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).

16. 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: 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.

Incorrect:

Not B, not E: Local and Global indexes can be automatically rebuild with UPDATE

INDEXES when you move the table.

17. Oracle Grid Infrastructure for a stand-alone server is installed on your production host before installing the Oracle Database server. The database and listener are configured by using Oracle Restart.

Examine the following command and its output:

```
$ crsctl config has
```

```
CRS-4622: Oracle High Availability Services auto start is enabled.
```

What does this imply?

- A. When you start an instance on a high with SQL *Plus dependent listeners and ASM disk groups are automatically started.
- B. When a database instance is started by using the SRVCTL utility and listener startup fails, the instance is still started.
- C. When a database is created by using SQL* Plus, it is automatically added to the Oracle Restart configuration.
- D. When you create a database service by modifying the SERVICE_NAMES initialization parameter, it is automatically added to the Oracle Restart configuration.

Answer: B

Explanation: About Startup Dependencies Oracle Restart ensures that Oracle components are started in the proper order, in accordance with component dependencies. For example, if database files are stored in Oracle ASM disk groups, then before starting the database instance, Oracle Restart ensures that the Oracle ASM instance is started and the required disk groups are mounted. Likewise, if a component must be shut down, Oracle Restart ensures that dependent components are cleanly shut down first. Oracle Restart also manages the weak dependency between database instances and the Oracle Net listener (the listener): When a database instance is started, Oracle Restart attempts to start the listener. If the listener startup fails, then the database is still started. If the listener later fails, Oracle Restart does not shut down and restart any database instances.

http://docs.oracle.com/cd/E16655_01/server.121/e17636/restart.htm#ADMIN12710

18. Which statement is true regarding the startup of a database instance?

- A. The instance does not start up normally and requires manual media recovery after a shutdown using the abort option.
- B. Uncommitted transactions are rolled back during the startup of the database instance after a shutdown using the immediate option.
- C. There is no difference in the underlying mechanics of the startup whether the database is shut down by using the immediate option or the abort option.
- D. Media recovery is required when the database is shut down by using either the immediate option or the abort option.
- E. Instance recovery is not required if the database instance was shut down by using SHUTDOWN IMMEDIATE.

Answer: B

Reference: http://docs.oracle.com/cd/A87860_01/doc/server.817/a76956/start.htm

19. Examine the memory-related parameters set in the SPFILE of an Oracle database:

memory_max_target=6G memory_target=5G pga_aggregate_target=500M sga_max_size=0 sga_target=0

Which statement is true?

- A. Only SGA components are sized automatically.
- B. Memory is dynamically re-allocated between the SGA and PGA as needed.
- C. The size of the PGA cannot grow automatically beyond 500 MB.
- D. The value of the MEMORY_TARGET parameter cannot be changed dynamically.

Answer: C

20. You create a new pluggable database, HR_PDB, from the seed database. Which three tablespaces are created by default in HR_PDB?

- A. SYSTEM
- B. SYSAUX
- C. EXAMPLE
- D. UNDO

E. TEMP

F. USERS

Answer: A,B,E

Explanation: * A PDB would have its SYSTEM, SYSAUX, TEMP tablespaces. It can also contain other user created tablespaces in it. *

* Oracle Database creates both the SYSTEM and SYSAUX tablespaces as part of every database.

* tablespace_datafile_clauses

Use these clauses to specify attributes for all data files comprising the SYSTEM and SYSAUX tablespaces in the seed PDB.

Incorrect:

Not D: a PDB can not have an undo tablespace. Instead, it uses the undo tablespace belonging to the CDB.

Note:

* Example: CONN pdb_admin@pdb1

```
SELECT tablespace_name FROM dba_tablespaces; TABLESPACE_NAME
```

```
SYSTEM SYSAUX TEMP USERS
```

```
SQL>
```

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 1Z0-062 Exam with Our Prep Materials Via below:

<https://www.certleader.com/1Z0-062-dumps.html>