DUMPSARENA

Oracle Database 12c: Installation and Administration

Oracle 1z0-062

Version Demo

Total Demo Questions: 20

Total Premium Questions: 411

Buy Premium PDF

https://dumpsarena.co

sales@dumpsarena.co

sales@dumpsarena.co
dumpsarena.co

QUESTION NO: 1

You are the DBA supporting an Oracle 11g Release 2 database and wish to move a table containing several DATE, CHAR, VARCHAR2, and NUMBER data types, and the table's indexes, to another tablespace.

The table does not have a primary key and is used by an OLTP application.

Which technique will move the table and indexes while maintaining the highest level of availability to the application?

A. Oracle Data Pump.

B. An ALTER TABLE MOVE to move the table and ALTER INDEX REBUILD to move the indexes.

C. An ALTER TABLE MOVE to move the table and ALTER INDEX REBUILD ONLINE to move the indexes.

D. Online Table Redefinition.

E. Edition-Based Table Redefinition.

ANSWER: D

Explanation:

* Oracle Database provides a mechanism to make table structure modifications without significantly affecting the availability of the table. The mechanism is called online table redefinition.Redefining tables online provides a substantial increase in availability compared to traditional methods of redefining tables.

* To redefine a table online:

Choose the redefinition method: by key or by rowid

* By key—Select a primary key or pseudo-primary key to use for the redefinition. Pseudo-primary keys are unique keys with all component columns having NOT NULL constraints. For thismethod, the versions of the tables before and after redefinition should have the same primary key columns. This is the preferred and default method of redefinition.

* By rowid—Use this method if no key is available. In this method, a hidden column named M_ROW$$ is added to the post-redefined version of the table. It is recommended that this column bedropped or marked as unused after the redefinition is complete. If COMPATIBLE is set to 10.2.0 or higher, the final phase of redefinition automatically sets this column unused. You can then use the ALTER TABLE ... DROP UNUSED COLUMNS statement to drop it.

You cannot use this method on index-organized tables.

Note:

* When you rebuild an index, you use an existing index as the data source. Creating an index in this manner enables you to change storage characteristics or move to a new tablespace.

Rebuilding an index based on an existing data source removes intra-block fragmentation. Compared to dropping the index and using the CREATE INDEX statement, re-creating an existing index offers better performance.

Incorrect:

Not E: Edition-based redefinition enables you to upgrade the database component of an application while it is in use, thereby minimizing or eliminating down time.

QUESTION NO: 2

Examine the following command:

CREATE TABLE (prod_id number(4),

Prod_name varchar2 (20),

Category_id number(30),

Quantity_on_hand number (3) INVISIBLE);

Which three statements are true about using an invisible column in the PRODUCTS table? (Choose three.)

A. The %ROWTYPE attribute declarations in PL/SQL to access a row will not display the invisible column in the output.

B. The DESCRIBE commands in SQL *Plus will not display the invisible column in the output.

C. Referential integrity constraint cannot be set on the invisible column.

D. The invisible column cannot be made visible and can only be marked as unused.

E. A primary key constraint can be added on the invisible column.

ANSWER: A B E

Explanation:

AB: You can make individual table columns invisible. Any generic access of a table does not show the invisible columns in the table. For example, the following operations do not display invisible columns in the output:

* SELECT * FROM statements in SQL

* DESCRIBE commands in SQL*Plus

* %ROWTYPE attribute declarations in PL/SQL* Describes in Oracle Call Interface (OCI)

Incorrect:

Not D: You can make invisible columns visible.

You can make a column invisible during table creation or when you add a column to a table, and you can later alter the table to make the same column visible.

QUESTION NO: 3

Your database supports an online transaction processing (OLTP) workload in which one of the applications creates a temporary table for a session and performs transactions on it. This consumes a lot of undo tablespace and generates lots of redo. Which two actions would you take to solve this problem? (Choose two.)

A. Increase the size of the temporary tablespace.

B. Enable Automatic Memory Management (AMM).

C. Enable undo retention guarantee.

D. Enable temporary undo for the database.

E. Increase the size of the redo log buffer.

ANSWER: A D

QUESTION NO: 4

You are about to plug a multi-terabyte non-CDB into an existing multitenant container database (CDB).

The characteristics of the non-CDB are as follows:

– Version: Oracle Database 11g Release 2 (11.2.0.2.0) 64-bit

– Character set: AL32UTF8

– National character set: AL16UTF16

– O/S: Oracle Linux 6 64-bit

The characteristics of the CDB are as follows:

– Version: Oracle Database 12c Release 1 64-bit

– Character Set: AL32UTF8

– National character set: AL16UTF16– O/S: Oracle Linux 6 64-bit

Which technique should you use to minimize down time while plugging this non-CDB into the CDB?

A. Transportable database

B. Transportable tablespace

C. Data Pump full export/import

D. The DBMS_PDB package

E. RMAN

ANSWER: B

Explanation:

* Overview, example:

- Log into ncdb12c as sys

- Get the database in a consistent state by shutting it down cleanly.

- Open the database in read only mode

- Run DBMS_PDB.DESCRIBE to create an XML file describing the database.

- Shut down ncdb12c

- Connect to target CDB (CDB2)

- Check whether non-cdb (NCDB12c) can be plugged into CDB(CDB2) - Plug-in Non-CDB (NCDB12c) as PDB(NCDB12c) into target CDB(CDB2).

- Access the PDB and run the noncdb_to_pdb.sql script.- Open the new PDB in read/write mode.

* You can easily plug an Oracle Database 12c non-CDB into a CDB. Just create a PDB manifest file for the non-CDB, and then use the manifest file to create a cloned PDB in the CDB. * Note that to plug in a non-CDB database into a CDB, the non-CDB database needs to be of version 12c as well. So existing 11g databases will need to be upgraded to 12c before they can be part of a 12c CDB.

QUESTION NO: 5

You plan to migrate your database from a File system to Automata Storage Management (ASM) on same platform.

Which two methods or commands would you use to accomplish this task? (Choose two.)

A. RMAN CONVERT command

B. Data Pump Export and import

C. Conventional Export and Import

D. The BACKUP AS COPY DATABASE . . . command of RMAN

E. DBMS_FILE_TRANSFER with transportable tablespace

ANSWER: A D

Explanation:

A:

1. Get the list of all datafiles.

Note: RMAN Backup of ASM Storage

There is often a need to move the files from the file system to the ASM storage and vice versa. This may come in handy when one of the file systems is corrupted by some means and then the file may need to be moved to the other file system.

D: Migrating a Database into ASM

* To take advantage of Automatic Storage Management with an existing database you must migrate that database into ASM. This migration is performed using Recovery Manager (RMAN) evenif you are not using RMAN for your primary backup and recovery strategy.

* Example:

Back up your database files as copies to the ASM disk group.

BACKUP AS COPY INCREMENTAL LEVEL 0 DATABASE

FORMAT '+DISK' TAG 'ORA_ASM_MIGRATION';

QUESTION NO: 6

You execute this command:

```
SQL> CREATE TABLESPACE lmtbsb DATAFILE '/u02/oracle/data/lmtbsb01.dbf' SIZE 50M
     EXTENT MANAGEMENT LOCAL;
```

Which two statements are true about segment space management for segments in this tablespace? (Choose two.)

A. Space utilization inside segments is mapped by bitmaps.

B. Segments are automatically shrunk and compressed when rows are deleted.

C. The PCTFREE storage parameter has no effect on segments created in this tablespace.

D. The PCTUSED storage parameter has no effect on segments created in this tablespace.

ANSWER: A D

QUESTION NO: 7

Identify three situations in which messages are written to the alert log file. (Choose three.)

A. Rebuilding an index using ALTER INDEX ... REBUILD fails with an error "ORA-01578: ORACLE data block corrupted (file # 14, block # 50)"

B. Creating a table returns "ORA-00955: name is already in used by an existing object"

C. Inserting a value into a table returns "ORA-01722: invalid number"

D. Updating a record in a table returns "ORA-00060: deadlock detected while waiting for resource"

E. Inserting a value into a table returns "ORA-00001: unique constraint (SYS.PK_XXXX) violated"

F. Running a query on a table returns "ORA-01578: ORACLE data block corrupted (file # 4, block # 131)"

ANSWER: A D F

QUESTION NO: 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 newlycreated database is a non-CDB and can never contain PDBs.

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

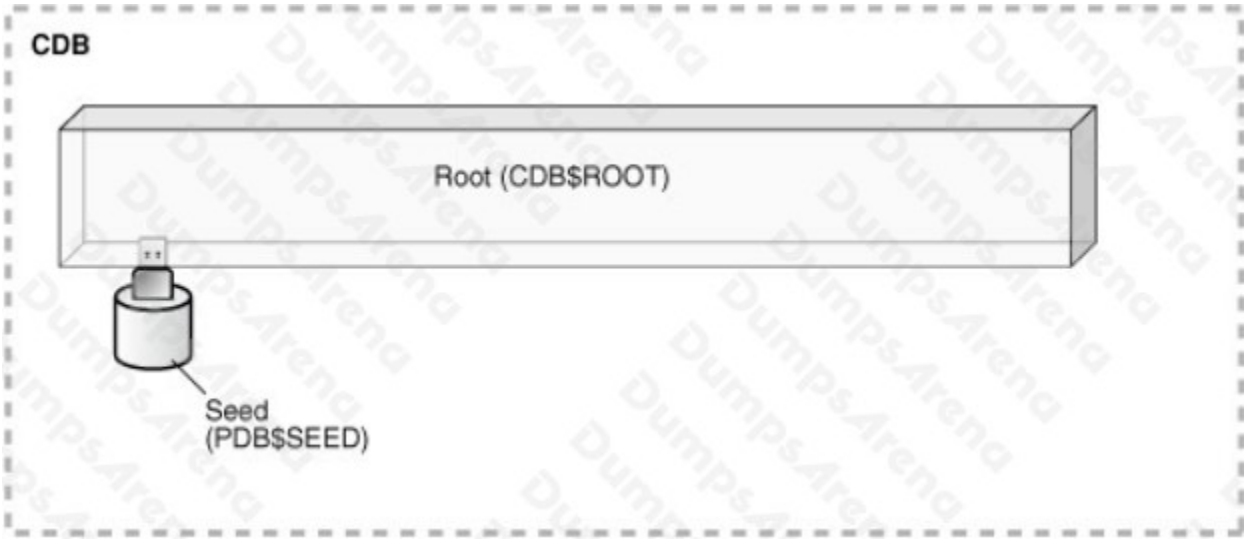

* Creating a PDB

Rather than constructing the data dictionary tables that define an empty PDB from scratch, and then populating its Obj$ and Dependency$ tables, the empty PDB is created when the CDB is created. (Here, we use empty to mean containing no customer-created artifacts.) It is referred to as the seed PDB and has the name PDB$Seed. Every CDB non-negotiably contains a seed PDB; it is non-negotiably always open in read-only mode. This has no conceptual significance; rather, it is just an optimization device. The create PDB operation is implemented as a special case of the clone PDB operation.

QUESTION NO: 9

Which three statements are true about Oracle checkpoint processing? (Choose three.)

A. Frequent thread checkpoints can degrade database performance

B. Database Writer (DBWn) processes write checkpoint information to datafile headers and the control file

C. It reduces the recovery time from instance failures

D. Incremental checkpoints write some dirty buffers to the datafiles and unwritten redo to the online redo logs.

E. Thread checkpoints ensure that all dirty buffers are written to data files during a normal shutdown

ANSWER: B C E

QUESTION NO: 10

Which two statements are true about Oracle network connections? (Choose two.)

A. A listener may listen on behalf of only one database instance at a time.

B. The listener check's a user's authentication credentials and creates a session if the credentials are valid.

C. The listener continuously monitors a connection after the user process connects to a service handler.

D. The listener can spawn a new server process to deal with each new connection.

E. A connection request from a remote client is always first received by a listener listening on port 1521.

ANSWER: B E

QUESTION NO: 11

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: B

QUESTION NO: 12

Examine the query and its output:

```
SQL> SELECT reason, metric_value FROM dba_outstanding_alerts;

REASON                                    METRIC_VALUE
----------------------------------        ------------------
Tablespace [TEST] is                      28.125
[28 percent] full

Metrics "Current Logons Count"            29
is at 29

Metrics "Database Time Spent              99.0375405
Waiting (%)" is at 99.03754 for
event class "Application"

db_recovery_file_dest_size of             97
4294967296 bytes is 97.298 used
and has 116228096 remaining
bytes available.
```

After 30 minutes, you execute the same query:

```
SQL> SELECT reason,metric_value FROM dba_outstanding_alerets;

REASON                                    METRIC_VALUE
----------------------------------        --------------
Tablespace [TEST] is                      28.125
[28 percent full
```

What might have caused three of the alerts to disappear?

A. The threshold alerts were cleared and transferred to DBA_ALERT_HISTORY.

B. An Automatic Workload Repository (AWR) snapshot was taken before the execution of the second query.

C. An Automatic Database Diagnostic Monitor (ADOM) report was generated before the execution of the second query.

D. The database instance was restarted before the execution of the second query.

ANSWER: D

QUESTION NO: 13

You administer an online transaction processing (OLTP) system whose database is stored in Automatic Storage Management (ASM) and whose disk group use normal redundancy.

One of the ASM disks goes offline, and is then dropped because it was not brought online before DISK_REPAIR_TIME elapsed.

When the disk is replaced and added back to the disk group, the ensuing rebalance operation is too slow.

Which two recommendations should you make to speed up the rebalance operation if this type of failure happens again? (Choose two.)

A. Increase the value of the ASM_POWER_LIMIT parameter.

B. Set the DISK_REPAIR_TIME disk attribute to a lower value.

C. Specify the statement that adds the disk back to the disk group.

D. Increase the number of ASMB processes.

E. Increase the number of DBWR_IO_SLAVES in the ASM instance.

ANSWER: A D

Explanation:

A: ASM_POWER_LIMIT specifies the maximum power on an Automatic Storage Management instance for disk rebalancing. The higher the limit, the faster rebalancing will complete. Lower values will take longer, but consume fewer processing and I/O resources.

D:

* Normally a separate process is fired up to do that rebalance. This will take a certain amount of time. If you want it to happen faster, fire up more processes. You tell ASM it can add moreprocesses by increasing the rebalance power.

* ASMB

ASM Background Process

Communicates with the ASM instance, managing storage and providing statistics

Incorrect:

Not B: A higher, not a lower, value of DISK_REPAIR_TIME would be helpful here.

Not E: If you implement database writer I/O slaves by setting the DBWR_IO_SLAVES parameter, you configure a single (master) DBWR process that has slave processes that are subservient to it. In addition, I/O slaves can be used to "simulate" asynchronous I/O on platforms that do not support asynchronous I/O or implement it inefficiently. Database I/O slaves provide non-blocking, asynchronous requests to simulate asynchronous I/O.

QUESTION NO: 14

Which statement is true about profiles?

A. Resource limits specified in a profile assigned to a user are always enabled.

B. A user can exist without any profile.

C. A profile can be assigned only to one user.

D. Password management using profiles is always enabled.

ANSWER: D

QUESTION NO: 15

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? (Choose two.)

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;

QUESTION NO: 16

The HR schema exists in databases, BOSTON and DENVER, with the same password, HR.

You have CREATE DATABASE LINK and CREATE SESSION privileges in both databases.

DENVER is defined as a service name in the tnsnames.ora of both databases.

You plan to use this command:

CREATE DATABASE LINK hr_link CONNECT TO hr IDENTIFIED BY hr USING 'DENVER';

What must be done to ensure that all users in BOSTON can access the HR schema in DENVER?

A. Change the command to create a public database link in BOSTON.

B. Change the command to create a public database link in DENVER.

C. Execute the command as SYS in DENVER.

D. Execute the command as SYS in BOSTON.

E. Execute the command as HR in BOSTON and SYS in DENVER.

F. Execute the command as SYS in both databases.

ANSWER: E

QUESTION NO: 17

You want to create a test database as a replica of your production database with minimum intervention from a DBA.

Which method would you use?

A. Use DBCA to create a template from the existing database to contain the database structure and then manually copy the data by using Oracle Data Pump.

B. Use Database Configuration Assistant (DBCA) to create a template from the existing database to contain the database structure.

C. Create the database by using the CREATE DATABASE... command and manually import data by using Data Pump.

D. Use DBCA to create a template from the existing database to contain the database structure with data files and then use the same template to create the database in the new location.

ANSWER: A

QUESTION NO: 18

Which set of statements is true about data dictionary views?

1. They are stored in the SYSTEM tablespace.

2. They are based on virtual tables.

3. They are owned by the SYS user.

4. They can be queried by a normal user only if the 07_DICTIONARY_ACCESSIBILITY parameter is set to TRUE.

5. The V$FIXED_TABLE view can be queried to list the names of these views.

6. They are owned by the SYSTEM user.

A. 2, 5, and 6

B. 1, 2, and 3

C. 1 and 3

D. 2, 3, 4, and 5

ANSWER: C

Explanation:

References: https://docs.oracle.com/cd/E11882_01/server.112/e40540/datadict.htm#CNCPT002

QUESTION NO: 19

The user SCOTT owns the CUST table that is placed in the SALES tablespace. The user SCOTT opens a session and executes commands as follows:

SQL> INSERT INTO cust VALUES(101, 'JACK');

1 row created.

SQL> INSERT INTO cust VALUES(102, 'SMITH'); 1 row created.

As a DBA, you execute the following command from another session: ALTER TABLESPACE sales READ ONLY;

Which statement is true regarding the effect of this command on the transaction in Scott's session?

A. The command fails as a transaction is still pending.

B. The transaction in Scott's session is rolled back and the tablespace becomes readonly.

C. The command waits and the user SCOTT can execute data manipulation language (DML) statements only as part of the current transaction.

D. The command hangs until all transactions on the objects in the tablespace commit or rollback, and then the tablespace is placed in readonly mode.

ANSWER: B

QUESTION NO: 20

Which two statements are true about availability audit features after migration to unified auditing? (Choose two.)

A. The ability of users to audit their own schema objects is not available in the post-migrated database.

B. Operating system audit trail is available in the post-migrated database.

C. Network auditing is available in the post-migrated database.

D. Mandatory auditing of audit administrative actions is available in the post-migrated database.

ANSWER: A D

Explanation:

References: https://docs.oracle.com/database/121/DBSEG/audit_changes.htm#DBSEG341