

2014 Latest Oracle 1Z0-060 Dumps Download With PDF & VCE (11-20)

QUESTION 11

Your multitenant container database (CDB) contains three pluggable database (PDBs). You find that the control file is damaged. You plan to use RMAN to recover the control file. There are no startup triggers associated with the PDBs. Which three steps should you perform to recover the control file and make the database fully operational?

- A. Mount the container database (CDB) and restore the control file from the control file auto backup.
- B. Recover and open the CDB in NORMAL mode.
- C. Mount the CDB and then recover and open the database, with the RESETLOGS option.
- D. Open all the pluggable databases.
- E. Recover each pluggable database.
- F. Start the database instance in the nomount stage and restore the control file from control file auto backup.

Answer: CDF

Explanation:

Step 1: F

Step 2: D

Step 3: C: If all copies of the current control file are lost or damaged, then you must restore and mount a backup control file. You must then run the RECOVER command, even if no data files have been restored, and open the database with the RESETLOGS option.

Note:

*RMAN and Oracle Enterprise Manager Cloud Control (Cloud Control) provide full support for backup and recovery in a multitenant environment. You can back up and recover a whole multitenant container database (CDB), root only, or one or more pluggable databases (PDBs).

QUESTION 12

A new report process containing a complex query is written, with high impact on the database. You want to collect basic statistics about query, such as the level of parallelism, total database time, and the number of I/O requests.

For the database instance STATISTICS_LEVEL, the initialization parameter is set to TYPICAL and the CONTROL_MANAGEMENT_PACK_ACCESS parameter is set to DIAGNOSTIC+TUNING.

What should you do to accomplish this task?

- A. Execute the query and view Active Session History (ASH) for information about the query.
- B. Enable SQL trace for the query.
- C. Create a database operation, execute the query, and use the DBMS_SQL_MONITOR.REPORT_SQL_MONITOR function to view the report.
- D. Use the DBMS_APPLICATION_INFO.SET_SESSION_LONGOPS procedure to monitor query execution and view the information from the V\$SESSION_LONGOPS view.

Answer: C

Explanation:

The REPORT_SQL_MONITOR function is used to return a SQL monitoring report for a specific SQL statement.

Incorrect:

Not A: Not interested in session statistics, only in statistics for the particular SQL query.

Not B: We are interested in statistics, not tracing.

Not D: SET_SESSION_LONGOPS Procedure

This procedure sets a row in the V\$SESSION_LONGOPS view. This is a view that is used to indicate the on-going progress of a long running operation. Some Oracle functions, such as parallel execution and Server Managed Recovery, use rows in this view to

indicate the status of, for example, a database backup.

Applications may use the SET_SESSION_LONGOPS procedure to advertise information on the progress of application specific long running tasks so that the progress can be monitored by way of the V\$SESSION_LONGOPS view.

QUESTION 13

Identify three valid options for adding a pluggable database (PDB) to an existing multitenant container database (CDB).

- A. Use the CREATE PLUGGABLE DATABASE statement to create a PDB using the files from the SEED.
- B. Use the CREATE DATABASE . . . ENABLE PLUGGABLE DATABASE statement to provision a PDB by copying file from the SEED.
- C. Use the DBMS_PDB package to clone an existing PDB.
- D. Use the DBMS_PDB package to plug an Oracle 12c non-CDB database into an existing CDB.
- E. Use the DBMS_PDB package to plug an Oracle 11 g Release 2 (11.2.0.3.0) non-CDB database into an existing CDB.

Answer: ACD

Explanation:

Use the CREATE PLUGGABLE DATABASE statement to create a pluggable database (PDB).

This statement enables you to perform the following tasks:

* (A) Create a PDB by using the seed as a template

Use the create_pdb_from_seed clause to create a PDB by using the seed in the multitenant container database (CDB) as a template. The files associated with the seed are copied to a new location and the copied files are then associated with the new PDB.

* (C) Create a PDB by cloning an existing PDB

Use the create_pdb_clone clause to create a PDB by copying an existing PDB (the source PDB) and then plugging the copy into the CDB. The files associated with the source PDB are copied to a new location and the copied files are associated with the new PDB.

This operation is called cloning a PDB.

The source PDB can be plugged in or unplugged. If plugged in, then the source PDB can be in the same CDB or in a remote CDB. If the source PDB is in a remote CDB, then a database link is used to connect to the remote CDB and copy the files.

* Create a PDB by plugging an unplugged PDB or a non-CDB into a CDB

Use the create_pdb_from_xml clause to plug an unplugged PDB or a non-CDB into a CDB, using an XML metadata file.

QUESTION 14

Your database supports a DSS workload that involves the execution of complex queries:

Currently, the library cache contains the ideal workload for analysis. You want to analyze some of the queries for an application that are cached in the library cache.

What must you do to receive recommendations about the efficient use of indexes and materialized views to improve query performance?

- A. Create a SQL Tuning Set (STS) that contains the queries cached in the library cache and run the SQL Tuning Advisor (STA) on the workload captured in the STS.
- B. Run the Automatic Workload Repository Monitor (AWRM).
- C. Create an STS that contains the queries cached in the library cache and run the SQL Performance Analyzer (SPA) on the workload captured in the STS.
- D. Create an STS that contains the queries cached in the library cache and run the SQL Access Advisor on the workload captured in the STS.

Answer: D

Explanation:

* SQL Access Advisor is primarily responsible for making schema modification recommendations, such as adding or dropping indexes and materialized views. SQL Tuning Advisor makes other types of recommendations, such as creating SQL profiles and

restructuring SQL statements.

*The query optimizer can also help you tune SQL statements. By using SQL Tuning Advisor and SQL Access Advisor, you can invoke the query optimizer in advisory mode to examine a SQL statement or set of statements and determine how to improve their efficiency. SQL Tuning Advisor and SQL Access Advisor can make various recommendations, such as creating SQL profiles, restructuring SQL statements, creating additional indexes or materialized views, and refreshing optimizer statistics.

Note:

* Decision support system (DSS) workload

*The library cache is a shared pool memory structure that stores executable SQL and PL/SQL code. This cache contains the shared SQL and PL/SQL areas and control structures such as locks and library cache handles.

QUESTION 15

The following parameter are set for your Oracle 12c database instance:

OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES=FALSE

OPTIMIZER_USE_SQL_PLAN_BASELINES=TRUE

You want to manage the SQL plan evolution task manually. Examine the following steps:

1. Set the evolve task parameters.
2. Create the evolve task by using the DBMS_SPM.CREATE_EVOLVE_TASK function.
3. Implement the recommendations in the task by using the DBMS_SPM.IMPLEMENT_EVOLVE_TASK function.
4. Execute the evolve task by using the DBMS_SPM.EXECUTE_EVOLVE_TASK function.
5. Report the task outcome by using the DBMS_SPM.REPORT_EVOLVE_TASK function.


Identify the correct sequence of steps:

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

Answer: B

Explanation:

*Evolving SQL Plan Baselines



2.Create the evolve task by using the DBMS_SPM.CREATE_EVOLVE_TASK function. This function creates an advisor task to prepare the plan evolution of one or more plans for a specified SQL statement. The input parameters can be a SQL handle, plan name or a list of plan names, time limit, task name, and description.

1.Set the evolve task parameters.

SET_EVOLVE_TASK_PARAMETER

This function updates the value of an evolve task parameter. In this release, the only valid parameter is TIME_LIMIT.

4. Execute the evolve task by using the DBMS_SPM.EXECUTE_EVOLVE_TASK function. This function executes an evolution task. The input parameters can be the task name, execution name, and execution description. If not specified, the advisor generates the name, which is returned by the function.

3: IMPLEMENT_EVOLVE_TASK

This function implements all recommendations for an evolve task. Essentially, this function is equivalent to using ACCEPT_SQL_PLAN_BASELINE for all recommended plans. Input parameters include task name, plan name, owner name, and execution name. 5. Report the task outcome by using the DBMS_SPM_EVOLVE_TASK function. This function displays the results of an evolve task as a CLOB. Input parameters include the task name and section of the report to include.

QUESTION 16

In a recent Automatic Workload Repository (AWR) report for your database, you notice a high number of buffer busy waits. The database consists of locally managed tablespaces with free list managed segments.

On further investigation, you find that buffer busy waits is caused by contention on data blocks.

Which option would you consider first to decrease the wait event immediately?

- A. Decreasing PCTUSED
- B. Decreasing PCTFREE
- C. Increasing the number of DBWnprocess
- D. Using Automatic Segment Space Management (ASSM)
- E. Increasing db_buffer_cache based on the V\$DB_CACHE_ADVICE recommendation

Answer: D

Explanation:

*Automatic segment space management (ASSM) is a simpler and more efficient way of managing space within a segment. It completely eliminates any need to specify and tune the pctused, freelists, and freelist groups storage parameters for schema objects created in the tablespace. If any of these attributes are specified, they are ignored. *Oracle introduced Automatic Segment Storage Management (ASSM) as a replacement for traditional freelists management which used one-way linked-lists to manage free blocks with tables and indexes. ASSM is commonly called "bitmap freelists" because that is how Oracle implements the internal data structures for free block management.

Note:

* Buffer busy waits are most commonly associated with segment header contention inside the data buffer pool (db_cache_size, etc.).

*The most common remedies for high buffer busy waits include database writer (DBWR) contention tuning, adding freelists (or ASSM), and adding missing indexes.

QUESTION 17

Examine this command:

```
SQL > exec DBMS_STATS.SET_TABLE_PREFS ('SH', 'CUSTOMERS', 'PUBLISH', 'false');
```

Which three statements are true about the effect of this command?

- A. Statistics collection is not done for the CUSTOMERS table when schema stats are gathered.
- B. Statistics collection is not done for the CUSTOMERS table when database stats are gathered.
- C. Any existing statistics for the CUSTOMERS table are still available to the optimizer at parse time.
- D. Statistics gathered on the CUSTOMERS table when schema stats are gathered are stored as pending statistics.
- E. Statistics gathered on the CUSTOMERS table when database stats are gathered are stored as pending statistics.

Answer: ACD

Explanation:

*SET_TABLE_PREFS Procedure

This procedure is used to set the statistics preferences of the specified table in the specified schema.

***Example:**

Using Pending Statistics

Assume many modifications have been made to the employees table since the last time statistics were gathered. To ensure that the cost-based optimizer is still picking the best plan, statistics should be gathered once again; however, the user is concerned that new statistics will cause the optimizer to choose bad plans when the current ones are acceptable. The user can do the following:

EXEC DBMS_STATS.SET_TABLE_PREFS('hr', 'employees', 'PUBLISH', 'false');

By setting the employees tables publish preference to FALSE, any statistics gather from now on will not be automatically published. The newly gathered statistics will be marked as pending.

QUESTION 18

Examine the following impdp command to import a database over the network from a pre-12c Oracle database (source):

```
$> impdp <user_name> full=Y network_link=hrdb_test transportable=always
transport_datafiles=
'/u01/app/oracle/oradata/hrdb/sales01.dbf',
'/u01/app/oracle/oradata/hrdb/cust01.dbf',
'/u01/app/oracle/oradata/hrdb/emp01.dbf'
version=12 logfile=import.log
```

Which three are prerequisites for successful execution of the command?

- A. The import operation must be performed by a user on the target database with the DATAPUMP_IMP_FULL_DATABASE role, and the database link must connect to a user on the source database with the DATAPUMP_EXP_FULL_DATABASE role.
- B. All the user-defined tablespaces must be in read-only mode on the source database.
- C. The export dump file must be created before starting the import on the target database.
- D. The source and target database must be running on the same platform with the same endianness.
- E. The path of data files on the target database must be the same as that on the source database.
- F. The impdp operation must be performed by the same user that performed the expdp operation.

Answer: BCE

Explanation:

* Be aware of the following general limitations as you plan to transport data:

/ The source and the target databases must use compatible database character sets.

C: You can use the full transportable export/import feature to copy an entire database from one Oracle Database instance to another. You can use Data Pump to produce an export dump file, transport the dump file to the target database if necessary, and then import the export dump file. Alternatively, you can use Data Pump to copy the database over the network.

B, Not D: Transporting a Database Using an Export Dump File The following list of tasks summarizes the process of transporting a database using an export dump file.

QUESTION 19

Which two are true concerning a multitenant container database with three pluggable database?

- A. All administration tasks must be done to a specific pluggable database.
- B. The pluggable databases increase patching time.
- C. The pluggable databases reduce administration effort.
- D. The pluggable databases are patched together.
- E. Pluggable databases are only used for database consolidation.

Answer: CDE

Explanation:

The benefits of Oracle Multitenant are brought by implementing a pure deployment choice. The following list calls out the most

compelling examples.

*High consolidation density.(E)

The many pluggable databases in a single multitenant 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. This is the same benefit that schema-based consolidation brings.

*Rapid provisioning and cloning using SQL.

*New paradigms for rapid patching and upgrades.(D, not B) The investment of time and effort to patch one multitenant container database results in patching all of its many pluggable databases. To patch a single pluggable database, you simply unplug/plug to a multitenant container database at a different Oracle Database software version.

* (C, not A)Manage many databases as one.

By consolidating existing databases as pluggable databases, administrators can manage many databases as one. For example, tasks like backup and disaster recovery are performed at the multitenant container database level.

*Dynamic between pluggable database resource management. In Oracle Database 12c, Resource Manager is extended with specific functionality to control the competition for resources between the pluggable databases within a multitenant container database.

Note:

*Oracle Multitenantis a new option for Oracle Database 12c Enterprise Edition that helps customers reduce IT costs by simplifying consolidation, provisioning, upgrades, and more. It is supported by a new architecture that allows a multitenant container database to hold many pluggable databases. And it fully complements other options, including Oracle Real Application Clusters and Oracle Active Data Guard. An existing database can be simply adopted, with no change, as a pluggable database; and no changes are needed in the other tiers of the application.

QUESTION 20

Examine the current value for the following parameters in your database instance:

SGA_MAX_SIZE = 1024M

SGA_TARGET = 700M

DB_8K_CACHE_SIZE = 124M

LOG_BUFFER = 200M

You issue the following command to increase the value of DB_8K_CACHE_SIZE:

```
SQL> ALTER SYSTEM SET DB_8K_CACHE_SIZE=140M;
```

Which statement is true?

- A. It fails because the DB_8K_CACHE_SIZE parameter cannot be changed dynamically.
- B. It succeeds only if memory is available from the autotuned components if SGA.
- C. It fails because an increase in DB_8K_CACHE_SIZE cannot be accommodated within SGA_TARGET.
- D. It fails because an increase in DB_8K_CACHE_SIZE cannot be accommodated within SGA_MAX_SIZE.

Answer: D

Explanation:

*The SGA_TARGET parameter can be dynamically increased up to the value specified for the SGA_MAX_SIZE parameter, and it can also be reduced.

* Example:

For example, suppose you have an environment with the following configuration:

SGA_MAX_SIZE = 1024M

SGA_TARGET = 512M

DB_8K_CACHE_SIZE = 128M

In this example, the value of SGA_TARGET can be resized up to 1024M and can also be reduced until one or more of the automatically sized components reaches its minimum size. The exact value depends on environmental factors such as the number of CPUs on the system. However, the value of DB_8K_CACHE_SIZE remains fixed at all times at 128M

*DB_8K_CACHE_SIZE

Size of cache for 8K buffers

* For example, consider this configuration:

SGA_TARGET = 512M

DB_8K_CACHE_SIZE = 128M

In this example, increasing DB_8K_CACHE_SIZE by 16 M to 144M means that the 16M is taken away from the automatically sized components. Likewise, reducing DB_8K_CACHE_SIZE by 16M to 112M means that the 16M is given to the automatically sized components.

If you want to pass the Oracle 12c 1Z0-060 exam successfully, recommend to read latest Oracle 12c 1Z0-060 Dumps full version.

<http://www.lead2pass.com/1z0-060.html>