1Z0-067 Q&As

Upgrade Oracle9i/10g/11g OCA to Oracle Database 12c OCP

Pass Oracle 1Z0-067 Exam with 100% Guarantee

Free Download Real Questions & Answers **PDF** and **VCE** file from:

https://www.passapply.com/1z0-067.html

100% Passing Guarantee
100% Money Back Assurance

Following Questions and Answers are all new published by Oracle Official Exam Center

QUESTION 1

Examine the command and its output:

SQL> DROP TABLE EMPLOYEE;

SQL> SELECT object_name AS recycle_name, original_name, type FROM recyclebin;

RECYCLE_NAMEORIGINAL_NAMETYPE

binsgk31sj/3akk5hg3j21kl5j3d==$0EMPLOYEE TABLE

You then successfully execute the command:

SQL> FLASHBACK TABLE "BINSgk31sj/3akk5hg3j21kl5j3d==$0" TO BEFORE DROP;

Which two statements are true?

A. It flashes back the employee table and all the constraints associated with the table.

B. It automatically flashes back all the indexes on the employes table.

C. It automatically flashes back any triggers defined on the table.

D. It flashes back only the structure of the table and not the data.

E. It flashes back the data from the recycle bin and the existing data in the original table is permanently lost.

Correct Answer: BC

A table and all of its dependent objects (indexes, LOB segments, nested tables, triggers, constraints and so on) go into the recycle bin together, when you drop the table. Likewise, when you perform Flashback Drop, the objects are generally all retrieved together. Trigger automatically flashbacked. References: http://docs.oracle.com/cd/B19306_01/backup.102/b14192/flashptr004.htm#i1020594

QUESTION 2

You want to reduce fragmentation and reclaim unused space for the sales table but not its dependent

objects. During this operation, you want to ensure the following:

i.Long-running queries are not affected.

ii.No extra space is used.

iii.Data manipulation language (DML) operations on the table succeed at all times throughout the process.

iv.Unused space is reclaimed both above and below the high water mark.

Which alter TABLE option would you recommend?

A. DEALLOCATE UNUSED

B. SHRINK SPACE CASCADE

C. SHRINK SPACE COMPACT

D. ROW STORE COMPRESS BASIC

Correct Answer: C

The COMPACT clause lets you divide the shrink segment operation into two phases. When you specify COMPACT, Oracle Database defragments the segment space and compacts the table rows but postpones the resetting of the high water mark and the deallocation of the space until a future time. This option is useful if you have long-running queries that might span the operation and attempt to read from blocks that have been reclaimed. The defragmentation and compaction results are saved to disk, so the data movement does not have to be redone during the second phase. You can reissue the SHRINK SPACE clause without the COMPACT clause during off-peak hours to complete the second phase. References: https://docs.oracle.com/cd/B28359_01/server.111/b28310/schema003.htm

QUESTION 3

Which two statements are true about scheduling operations in a pluggable database (PDB)?

A. Scheduler jobs for a PDB can be defined only at the container database (CDB) level.

B. A job defined in a PDB runs only if that PDB is open.

C. Scheduler attribute setting is performed only at the CDB level.

D. Scheduler objects created by users can be exported or imported using Data Pump.

E. Scheduler jobs for a PDB can be created only by common users.

Correct Answer: BD

In general, all scheduler objects created by the user can be exported/imported into the PDB using data pump. Predefined scheduler objects will not get exported and that means that any changes made to these objects by the user will have to be made once again after the database has been imported into the pluggable database. However, this is how import/export works currently. A job defined in a PDB will run only if a PDB is open.

QUESTION 4

Which two statements are true about roles in multitenant container databases (CDBs)?

A. Local roles can be granted to local and common users.

B. A common role can be granted only to a common user.

C. A common user can create a local role by default in any pluggable database (PDB) that is plugged in to a CDB.

D. A common role can be granted only system privileges.

E. The root container can have both local and common roles.

F. A local role can be assigned to a common role in a PDB.

Correct Answer: DF

QUESTION 5

You want to capture column group usage and gather extended statistics for better cardinality estimates for the customers table in the SH schema.

Examine the following steps:

1. Issue the SELECTDBMS_STATS. CREATE_EXTENDED_STATS(`SH\\', \\'CUSTOMERS\\')from dual

statement.

2.Execute the dbms_stats.seed_col_usage (null,`SH\\',500) procedure.

3.Execute the required queries on the customers table.

4.Issue the select dbms_stats.reportwcol_usage(`SH\\', \\'customers\\') from dual statement.

Identify the correct sequence of steps.

A. 3, 2, 1, 4

B. 2, 3, 4, 1

C. 4, 1, 3, 2

D. 3, 2, 4, 1

Correct Answer: B

Step 1 (2). Seed column usage Oracle must observe a representative workload, in order to determine the appropriate column groups. Using the new procedure DBMS_STATS.SEED_COL_USAGE, you tell Oracle how long it should observe the workload. Step 2: (3) You don\\'t need to execute all of the queries in your work during this window. You can simply run explain plan for some of your longer running queries to ensure column group information is recorded for these queries. Step 3. (1) Create the column groups At this point you can get Oracle to automatically create the column groups for each of the tables based on the usage information captured during the monitoring window. You simply have to call the DBMS_STATS.CREATE_EXTENDED_STATS function for each table.This function requires just two arguments, the schema name and the table name. From then on, statistics will be maintained for each column group whenever statistics are gathered on the table.

Note:

*

DBMS_STATS.REPORT_COL_USAGE reports column usage information and records all the SQL

operations the database has processed for a given object.

*

The Oracle SQL optimizer has always been ignorant of the implied relationships between data columns

within the same table. While the optimizer has traditionally analyzed the distribution of values within a

column, he does not collect value- based relationships between columns.

*

Creating extended statistics

Here are the steps to create extended statistics for related table columns

withdbms_stats.created_extended_stats:

1 - The first step is to create column histograms for the related columns.

2 ?Next, we run dbms_stats.create_extended_stats to relate the columns together.

Unlike a traditional procedure that is invoked via an execute ("exec") statement, Oracle extended statistics

are created via a select statement.