

Exam : Oracle 1Z0-054

**Title : Oracle Database 11g:
Performance Tuning**

Version : Demo

1. Identify two correct statements about the Active Session History (ASH) data. (Choose two.)

- A. A part of SGA memory is used to store ASH data as rolling buffer.
- B. The ASH data can be analyzed between any two small time intervals.
- C. All ASH data in memory are flushed to disk by MMON in every 3 seconds.
- D. All ASH data in memory are flushed to disk by MMNL process whenever the buffer is full.

Answer: AB

2. A user in a session executed the following SQL statement to set the optimizer mode:

```
ALTER SESSION SET OPTIMIZER_MODE = ALL_ROWS
```

What impact would it have on the goal of the optimizer for that session? (Choose all that apply.)

- A. Statement level OPTIMIZER_MODE hints take precedence over the session-level setting.
- B. The OPTIMIZER_MODE parameter set at instance level takes precedence over the session-level value.
- C. The optimizer uses a cost-based approach, regardless of the presence of statistics; it optimizes with a goal of best response time.
- D. The optimizer uses a cost-based approach for all SQL statements in the session, regardless of the presence of statistics; it optimizes with a goal of best throughput.

Answer: AD

3. You are working on an online transaction processing (OLTP) system. By day most of the application users perform queries accessing the most recently added or modified rows. The applications have most of the queries based on multiple tables. But at night, some batch processing is also done.

Which two actions would you recommend to choose a goal for the optimizer based on the needs of your application? (Choose two.)

- A. setting the OPTIMIZER_MODE parameter to ALL_ROWS at the instance level
- B. setting the OPTIMIZER_MODE parameter to FIRST_ROWS_n at the instance level
- C. asking the developer to add a hint /*ALL_ROWS*/ in the long-running batch processing queries
- D. asking the developer to add a hint /*FIRST_ROWS_n*/ in the long-running batch processing queries

Answer: BC

4. Which three factors influence the optimizer's behavior while choosing an optimization approach and goal for a SQL statement? (Choose three.)

- A. parsing of a SQL statement
- B. operating system (OS) statistics
- C. object statistics in the data dictionary
- D. the OPTIMIZER_MODE initialization parameter
- E. optimizer SQL hints for changing the query optimizer goal

Answer: CDE

5. Examine the initialization parameter values for the instance given below:

name	TYPE	VALUE
optimizer_capture_sql_plan_baselines	boolean	FALSE
optimizer_dynamic_sampling	integer	2
optimizer_features_enable	string	11.1.0.6
optimizer_index_caching	integer	0
optimizer_index_cost_adj	integer	100
optimizer_mode	string	ALL_ROWS
db_file_multiblock_read_count	integer	64

You notice that the one of the queries is using a full table scan (view Exhibit1) instead of index unique scan (view Exhibit2). The index is present on the column that is accessed in the WHERE clause of the query. The cost for a full table scan is more than that for an index unique scan.

Why would the optimizer choose full table scan over index unique scan? (Choose all that apply.)

- A. The OPTIMIZER_INDEX_COST_ADJ initialization parameter is set to a low value.
- B. The OPTIMIZER_INDEX_COST_ADJ initialization parameter is set to a high value.
- C. The DB_FILE_MULTIBLOCK_READ_COUNT initialization parameter is set to a low value.
- D. The statistics for the table and all the indexes associated with the table are not current.

Answer: BD

6. The columns CUST_CITY, CUST_STATE_PROVINCE, and COUNTRY_ID are frequently used together in the WHERE clause of the queries. The CUSTOMERS table is a big table with 20 GB of data. You notice that the selectivity for these three columns varies from the selectivity that the optimizer calculates.

What would you recommend to influence the selectivity calculated by the optimizer?

- A. creating function-based indexes by concatenating all the columns
- B. updating histogram statistics for these columns by using the DBMS_STATS.GATHER_TABLE_STATS procedure
- C. using the DBMS_STATS.CREATE_EXTENDED_STATS function to create a virtual column and create index on the virtual column
- D. using the DBMS_STATS.CREATE_EXTENDED_STATS function to create a virtual column and the DBMS_STATS.GATHER_TABLE_STATS procedure to collect statistics on the virtual column

Answer: D

7. View Exhibit1 to examine the description of the CUSTOMERS table.

You observed that optimizer selectivity is not accurate when the CUST_STATE_PROVINCE and COUNTRY_ID columns are used together in the WHERE clause of a query.

View Exhibit2 to examine the query execution plan and the commands executed to gather the statistics.

The optimizer predicts that 20 rows will be processed rather than the 3,341 rows, which is the actual number of rows returned from the table.

What can you do to make the optimizer detect the actual number of rows?

- A. Set the STATISTICS_LEVEL parameter to ALL.
- B. Set the OPTIMIZER_USE_PENDING_STATISTICS parameter to FALSE.
- C. Create extended statistics for the CUST_STATE_PROVINCE and COUNTRY_ID columns.
- D. Increase the STALE_PERCENT value for the CUSTOMERS table by using the DBMS_STATS.SET_TABLE_PREFS procedure.

Answer: C

8. View Exhibit1 to examine the description of the CUSTOMERS table.

The CUSTOMERS table has been updated heavily today. In a frequently used SQL statement, you notice that estimated rows and the actual number of rows fetched differ greatly. The COUNTRY_ID column has an index.

View Exhibit2 and examine the query execution plan.

What would you recommend to improve the optimizer's estimation?

- A. setting the STATISTICS_LEVEL parameter to ALL
- B. setting the OPTIMIZER_USE_PENDING_STATISTICS parameter to FALSE
- C. creating extended statistics for the CUST_LAST_NAME, CUST_ID, and CUST_TOTAL columns
- D. updating the statistics for the CUSTOMERS table by using the DBMS_STATS.GATHER_TABLE_STATS procedure

Answer: D