



[Oracle 1z0-054](#)

Exam Name: Oracle Database 11g: Performance Tuning

Q & A : 192 Q&As

Pdf Demo

Quality and Value for the 1z0-054 Exam

[Just4Exams Practice Exams](#) for Oracle Oracle 11i 1z0-054 are written to the highest standards of technical accuracy, using only certified subject matter experts and published authors for development.

100% Guarantee to Pass Your 1z0-054 Exam

If you do not pass the Oracle 11i 1z0-054 exam on your first attempt using our Just4Exams **1z0-054 testing engine and pdf study guide**, we will give you a FULL REFUND of your purchasing fee.

Downloadable, Interactive 1z0-054 Testing engines and PDF Version

Our Exam Preparation Material provides you everything you will need to take a [Oracle 11i certification](#) examination. Details are researched and produced by [Oracle Certification](#) Experts who are constantly using industry experience to produce precise, and logical.

Free 1z0-054 Exams:

This is demo only, this pdf do not include the questions and answers picture

Exam : Oracle 1Z0-054

Title : Oracle Database 11g: Performance Tuning

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. View Exhibit1 and examine the indexes on the CUSTOMERS table.

The statistics for the CUSTOMERS table have been updated recently by using the following command:

SQL> EXEC DBMS_STATS.GATHER_TABLE_STATS('SH','CUSTOMERS',method_opt=>'FOR ALL INDEXED COLUMNS SIZE AUTO');
View Exhibit2 to examine a query plan. Even though the index is present on the COUNTRY_ID and CUST_GENDER columns, the query uses a full table scan. What could be the reason?

- A. because the histogram statistics for the COUNTRY_ID column are not updated
- B. because the DB_FILE_MULTIBLOCK_READ_COUNT initialization parameter is set to a high value
- C. because the optimizer calculates the cost of accessing blocks by using a full table scan to be less as compared to index scans, even though indexes are available
- D. because indexes on CUST_GENDER and COUNTRY_ID columns are of different types, the index on the CUST_GENDER column is bitmap index, and on COUNTRY_ID columns is btree index.

Answer: C

3. 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

4. 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

5. 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

6. 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

7. You are working on a decision support system (DSS). The index is available on the COUNTRY_ID column of the CUSTOMERS table.

View the Exhibit and examine the parameter settings and the query execution plan.

Why is the query using a full table scan instead of an index scan?

- A. because the histogram statistics for the COUNTRY_ID column are not updated
- B. because the index statistics for the index on the COUNTRY_ID column are not current

- C. because the DB_FILE_MULTIBLOCK_READ_COUNT initialization parameter is set to a higher value
- D. because the optimizer predicts that most of the blocks in the table are accessed. Therefore, it uses a full table scan, even though indexes are available.

Answer: D

8. View the Exhibit exhibit1 to examine the series of SQL commands and parameter settings.

View the Exhibit exhibit2 to examine the plans available in the SQL plan baseline.

The first plan (in red) is created when OPTIMIZER_MODE is set to ALL_ROWS and the second plan (in blue) is created when OPTIMIZER_MODE is set to FIRST_ROWS.

Which SQL plan baseline would be used if the SQL query in exhibit1 is executed again when the value of OPTIMIZER_MODE is set to FIRST_ROWS?

- A. the second plan, because it is a fixed plan
- B. the first plan, because it is an accepted plan
- C. the second plan, because it is the latest generated plan in FIRST_ROW mode
- D. A new plan, because the second plan in FIRST_ROW mode is not an accepted plan

Answer: B

9. 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

10. You are working on an online transaction processing (OLTP) system. You detected free buffer waits events for your database instance running in a machine that has multiple CPUs. You increased the database buffer cache size as the first step. After a few hours of work on the database, further investigation shows that the same event is being recorded.

What would be your next step to avoid this event in future?

- A. Decrease the value of the DBWR_IO_SLAVES parameter.
- B. Set the USE_INDIRECT_DATA_BUFFERS parameter to TRUE.
- C. Increase the value of the DB_WRITER_PROCESSES parameter.
- D. Increase the value of the DB_FILE_MULTIBLOCK_READ_COUNT parameter.

Answer: C

11. A batch workload that historically completed in the maintenance window between 10:00 PM and midnight is currently showing poor performance and completing at 2 AM.

To help in the diagnosis of the performance degradation, the senior DBA in your organization asks you to execute the awrdrrpt.sql script to produce a Compare Periods report.

Which two statements are true about the report produced by this script? (Choose two.)

- A. It is refreshed automatically based on the moving window baseline.
- B. It compares details between any two selected time periods of the same duration.
- C. It normalizes the statistics by the amount of time spent on the database for each time period.
- D. It compares details between two consecutive time periods of the same or different durations and is refreshed every 60 minutes.

Answer: BC

12. You are working on a development database that was upgraded to Oracle Database 11g from Oracle Database 9i. An ADDM finding in this database says that the shared pool is inadequately sized, as shown in the Exhibit.

You diagnosed that this is due to different kinds of workloads and this occurs only during peak hours. You tried to resize this by shrinking the database buffer cache but that caused inadequate buffer cache problems. The following are the related parameter settings:

```
SQL> show parameter sga
```

```
name TYPE VALUE
```

```
-----  
lock_sga boolean FALSE
```

```
pre_page_sga boolean FALSE
```

```
sga_max_size big integer 300M
```

```
sga_target big integer 0
```

```
SQL> show parameter target
```

```
name TYPE VALUE
```

.....
fast_start_mttr_target integer 0
memory_max_target big integer 0
memory_target big integer 0
pga_aggregate_target big integer 100M
sga_target big integer 0

You want to balance the memory between the System Global Area (SGA) components within SGA without affecting the size of the Program Global Area (PGA).

Which action would solve this problem?

- A. Set the SGA_TARGET parameter to 300M.
- B. Set the SGA_MAX_SIZE parameter to 400M.
- C. Set the MEMORY_TARGET parameter to 100M.
- D. Set the MEMORY_MAX_TARGET parameter to 300M.

Answer: A

13. 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

14. 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

The index created on the column used in the WHERE clause of the query. You notice that the query is not using the index. Instead of an index scan, a full table scan is used.

View the Exhibit and examine the autotrace output for a query.

What could be the reason for it? (Choose all that apply.)

- A. The OPTIMIZER_INDEX_COST_ADJ initialization parameter has a low value.
- B. The DB_FILE_MULTIBLOCK_READ_COUNT initialization parameter has a low value.
- C. The statistics for the table and all the indexes associated with the table are not current.
- D. The table has less than DB_FILE_MULTIBLOCK_READ_COUNT blocks under the high-water mark.

Answer: CD

15. After running SQL Performance Analyzer (SPA), you observe a few regressed SQL statements in the SPA output. Identify the two actions that you would suggest for these regressed SQL statements. (Choose two.)

- A. Running SQL Access Advisor
- B. Adding them to SQL Plan Baseline
- C. Submitting them to SQL Tuning Advisor

D. Running Automatic Database Diagnostic Monitor (ADDM)

Answer: BC

16. You observed that some of the queries are performing poorly on the SALES_RECORDS table in your database. On further investigation, you find that at the end of each day the contents of the SALES_RECORDS table are transferred to the SALES table and deleted from the SALES_RECORDS table. The deleted operations cause the table to be sparsely populated. You decided to use the ALTER TABLE...SHRINK SPACE COMPACT command to shrink the table.

Why would you choose this method? (Choose all that apply.)

- A. because it can be used during peak hours
- B. because it avoids unnecessary cursor invalidation
- C. because it adjusts the high-water mark (HWM) immediately
- D. because you have long-running queries that might span the shrink operation
- E. because it does not allow any data manipulation language (DML) operations, thereby making the shrink operation faster

Answer: ABD

17. 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

18. View the Exhibit and examine a portion of the output obtained from the following query:

```
SQL> SELECT * FROM v$sqlsys_time_model;
```

Select three correct interpretations of the time model statistics. (Choose three.)

- A. DB time includes the wait time of all the nonidle and idle user sessions.
- B. SQL execute elapsed time includes the time spent in performing fetches of query results.
- C. DB CPU includes the CPU time spent on database user-level calls and background CPU time.
- D. SQL execute elapsed time includes components of the hard parse elapsed time like bind elapsed time.
- E. DB time includes the connection management call elapsed time excluding the background process time.

Answer: BDE

19. You work as a DBA for a company and you have the responsibility of managing one of its online transaction processing (OLTP) systems. The database encountered performance-related problems and you generated an Automatic Workload Repository (AWR) report to investigate it further.

View the Exhibits and examine the AWR report.

What could be the problem in this database?

- A. Java pool is not configured.
- B. The CPU in the system is slow.
- C. The shared pool size is inadequate.
- D. The database buffer cache is inadequate.
- E. The OPEN_CURSORS parameter is set to a small value.

Answer: C

20. You plan to use adaptive thresholds as part of the performance tuning activity. You decide to increase the window size of the default moving window baseline for all metric observations and comparisons in your database. The following error occurs when you try to increase the window size through Enterprise Manager:

```
Failed to commit: ORA-13541: system moving window baseline size (1296000) greater than retention (1036800) ORA-06512: at "SYS.DBMS_WORKLOAD_REPOSITORY", line 601 ORA-06512: at line 2
```

Which action would allow you to perform the preceding task successfully?

- A. increasing the flashback retention period
- B. increasing the retention period for SQL Management Base
- C. increasing the undo retention period for the database instance
- D. increasing the Automatic Workload Repository (AWR) retention period

Answer: D

More [1z0-054 Braindumps](#) Information

1Z0-052	1z0-054	1Z0-200	1z0-050	1z0-236
---------	---------	---------	---------	---------

1z0-032	1Z0-110	1z0-141	1z0-271	1Z0-144
1z0-010	1Z1-402	1Z0-055	1z0-255	1z1-312
1Z0-114	1Z0-215	1Z1-054	1z0-095	1Z0-518
1Z0-146	1z1-047	1Z0-519	1z0-600	1z0-403