

# DUMPS ARENA

## Oracle Database 11g: Performance Tuning

Oracle 1z0-054

Version Demo

Total Demo Questions: 10

Total Premium Questions: 192

Buy Premium PDF

<https://dumpsarena.co>

[sales@dumpsarena.co](mailto:sales@dumpsarena.co)

[sales@dumpsarena.co](mailto:sales@dumpsarena.co)  
[dumpsarena.co](https://dumpsarena.co)

**QUESTION NO: 1**

You observe that suboptimal execution plans for the queries are being generated on a table that previously used less resources. You have collected statistics on these tables two days ago. The optimizer statistics retention period is set to 31 days. You are able to find the timestamp information about statistics update from the DBA\_TAB\_STATS\_HISTORY view. Because it is a frequently queried table, you would like the optimizer to generate better plans.

Which action would enable you to use the previous set of statistics on the objects that may lead to better execution plans?

- A. restoring statistics from statistics history up to the desired time
- B. deleting all AWR snapshots collected after the time of desired statistics collection
- C. applying the flashback table technique until the time of desired statistics collection
- D. setting the OPTIMIZER\_PENDING\_STATISTICS parameter to TRUE to use the previous version of statistics

**ANSWER: A****QUESTION NO: 2**

You are working on a 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. 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 the SGA depending on the workloads. Which two options would solve this problem? (Choose two.)

- A. Set the SGA\_TARGET parameter to 300M.
- B. Set the SGA\_MAX\_SIZE parameter to 400M.
- C. Set the MEMORY\_MAX\_TARGET and MEMORY\_TARGET parameters to 400M.
- D. Set the MEMORY\_MAX\_TARGET and MEMORY\_TARGET parameters to 100M.
- E. Set the PGA\_AGGREGATE\_TARGET parameter to 200M and the SGA\_MAX\_SIZE parameter to 400M.

**ANSWER: A C**

### QUESTION NO: 3

Examine the output of the following query:

```
SQL> SELECT c.name,a.addr,a.gets,a.misses,a.sleeps,
```

```
2 a.immediate_gets,a.immediate_misses,b.pid
```

```
3 FROM v$latch a, v$latchholder b, v$latchname c
```

```
4 WHERE a.addr = b.laddr(+) and a.latch# = c.latch#
```

```
5 AND c.name LIKE '&latch_name%' ORDER BY a.latch#;
```

```
LATCH NAME ADDR GETS MISSES SLEEPS IMMEDIATE_GETS IMMEDIATE_MISSES
```

-----

```
shared pool 20016544 8520540 14112 3137 0 0
```

Which two conclusions can you draw from this? (Choose two.)

- A. The latch was requested in no wait mode.

- B.** The latch was requested in willing-to-wait mode.
- C.** The shared pool operations to allocate and free memory in it caused the misses.
- D.** Automatic Memory Management is not enabled because the IMMEDIATE\_GETS and IMMEDIATE\_MISSES columns have zero in them.

**ANSWER: B C**

**QUESTION NO: 4**

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.

```

NAME                                TYPE                                VALUE
-----                                -
db_file_multiblock_read_count       integer                             49

SQL> SELECT BLOCKS, EMPTY_BLOCKS FROM ALL_TABLES
WHERE TABLE_NAME = 'CUST';

BLOCKS EMPTY_BLOCKS
-----
2902          0

SQL> EXEC DBMS_STATS.GATHER_TABLE_STATS('SH','CUST');

PL/SQL procedure successfully completed.

SQL> select cust_id, cust_last_name from sh.cust where country_id=52790;

74080 rows selected.

Execution Plan
-----
Plan hash value: 260468903

-----
| Id | Operation          | Name | Rows  | Bytes | Cost (%CPU)| Time     |
-----
|  0 | SELECT STATEMENT   |      | 74150 | 1303K | 793  (1)| 00:00:10 |
|*  1 | TABLE ACCESS FULL| CUST | 74150 | 1303K | 793  (1)| 00:00:10 |
-----

Predicate Information (identified by operation id):
-----

   1 - filter("COUNTRY_ID"=52790)

Statistics
-----
      0 recursive calls
      0 db block gets
    7659 consistent gets
    2848 physical reads
      0 redo size
  1584703 bytes sent via SQL*Net to client
    54738 bytes received via SQL*Net from client
    4940 SQL*Net roundtrips to/from client

```

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**

**QUESTION NO: 5**

You plan to collect statistics for an important table in your database to help improved query performance. You are uncertain about the optimizer behavior after the statistics have been collected because currently there are no statistics for indexes or partitions on the table. You wish to test the performance effect of the newly collected statistics. How would you prevent the optimizer from using the newly collected statistics for other sessions immediately after statistics collection?

- A. Set the PUBLISH preference to FALSE for the table.
- B. Set the NO\_VALIDATE preference to TRUE for the table.
- C. Set the OPTIMIZER\_DYNAMIC\_SAMPLING parameter to 0.
- D. Set the OPTIMIZER\_USE\_SQL\_PLAN\_BASELINES parameter to FALSE for the session.
- E. Set the OPTIMIZER\_USE\_PENDING\_STATISTICS parameter to FALSE for the session.

**ANSWER: A****QUESTION NO: 6**

Examine the output of the following query:

```
SQL> SELECT c.name,a.addr,a.gets,a.misses,a.sleeps,
```

```
2 a.immediate_gets,a.immediate_misses,b.pid
```

```
3 FROM v$latch a, v$latchholder b, v$latchname c
```

```
4 WHERE a.addr = b.laddr(+) and a.latch# = c.latch#
```

```
5 AND c.name LIKE '&latch_name%' ORDER BY a.latch#;
```

```
LATCH NAME ADDR GETS MISSES SLEEPS IMMEDIATE_GETS IMMEDIATE_MISSES
```

```
-----  
shared pool 20016544 8520540 14112 3137 0 0
```

You calculated the Gets-to-Misses ratio, which is .99834, and this ratio is dropping over a period of time.

Which two actions can improve this? (Choose two.)

- A. reducing hard parses
- B. increasing the size of the shared pool
- C. using only dedicated server connections
- D. setting the CURSOR\_SHARING parameter to EXACT
- E. encouraging the use of more literal SQL statements

**ANSWER: A B**

**QUESTION NO: 7**

You have the automatic optimizer statistics gathering enabled for a 24x7 database as part of the automatic maintenance task. Because the automatic optimizer statistics collection runs during maintenance window(in the night), the statistics on tables that are significantly modified throughout the day becomes stale. This leads to poor query performance on these tables and you want to minimize the

overhead in statistics collection. Identify a solution to this problem.

- A. Change the STALE\_PERCENT preference to zero for the tables.
- B. Set the OPTIMIZER\_USE\_PENDING parameter statistics to TRUE.
- C. Set the OPTIMIZER\_DYNAMIC\_SAMPLING parameter to zero during day time.
- D. Gather statistics on the table when the tables have a representative number of rows, and then lock the statistics.

**ANSWER: D****QUESTION NO: 8**

View the Exhibit and analyze the content of V\$BUFFER\_POOL\_STATISTICS.

```
SQL> ALTER TABLE scott.emp STORAGE(BUFFER_POOL KEEP);
```

```
Table altered.
```

```
SQL> SELECT name, buf_get FROM v$buffer_pool_statistics WHERE name='KEEP';
```

NAME	BUF_GET
KEEP	376

```
SQL> SELECT * FROM scott.emp;
```

EMPNO	ENAME	JOB	MGR	HIREDATE	SAL	COMM	DEPTNO
7369	SMITH	CLERK	7902	17-DEC-80	800		20
7499	ALLEN	SALESMAN	7698	20-FEB-81	1600	300	30
7521	WARD	SALESMAN	7698	22-FEB-81	1250	500	30

```
.....
.....
.....
```

```
14 rows selected.
```

```
SQL> SELECT name, buf_get FROM v$buffer_pool_statistics WHERE name='KEEP';
```

NAME	BUF_GET
KEEP	376

Why does the query on the EMP table not access buffers from the keep pool?

- A. because the query used SCOTT.EMP instead of EMP
- B. because the query with \* always uses the default pool
- C. because the query with \* always forces physical I/Os that bypass the keep pool

D. because the blocks of the EMP table are already available in one of the other buffer pools

**ANSWER: D**

### QUESTION NO: 9

You are working as a DBA in ABC Corp. You are working on online transaction processing (OLTP) system. The database uses ASM storage. One of the ASM disk goes offline because of hardware failure. When the disk is added back again after the fix, because the database is active, rebalance operations impact the I/O subsystem. Which two recommendations would you give to lower the impact on the I/O subsystem? (Choose two.)

- A. Increase the number of ASMB processes.
- B. Increase the number of DBWR\_IO\_SLAVES.
- C. Decrease the value for the ASM\_POWER\_LIMIT parameter.
- D. Increase the value for the ASM\_POWER\_LIMIT parameter.
- E. Set the DISK\_REPAIR\_TIME disk attribute to a lower value.
- F. Specify the POWER clause with a lower value in an ALTER DISKGROUP statement.

**ANSWER: C F**

### QUESTION NO: 10

You are working on an online transaction processing (OLTP) system. The middle-tier applications use connection pooling to connect to the database. Presently, you have a single-node database. The company plans to migrate the database to a RAC environment. Before you move to a RAC environment, you want to test the performance of the SQL statements and the peak workload on the new environment. To accomplish the Database Capture for replay, you identified the peak workload period on the existing system and started the Database Capture. Which client requests to the database can be captured as part of the workload capture? (Choose all that apply.)

- A. flashback query
- B. distributed transactions
- C. logging in and logging out of sessions
- D. all DDL statements having bind variables
- E. direct path load of data from external files

**ANSWER: C D**