

Oracle

Exam 1z0-054

Oracle Database 11g: Performance Tuning

Version: 5.0

[Total Questions: 192]

Question No : 1

You work for a small manufacturing company as a DBA. The company has various applications to manage the details of its business and customer base. The company has only one sales outlet where the operator updates the SALES table in the sales application with the details when a product is sold. The company has many other applications for various other businesses that use the same database.

The sales operator, while generating a report, noticed that there are some gaps generated in the serial number in the SLNO column during peak hours of business.

On investigating, you find that the SLNO column in the SALES table uses the SLNSEQ sequence to generate the serial numbers when a record is inserted.

During peak hours, you noticed the following:

When the operator entered the sales details the last time, the sequence number was 1056300, but when you checked the sequence as follows, you found a gap of 100 numbers:

```
SQL> SELECT sequence_name, last_number FROM user_sequences WHERE  
sequence_name='SLNSEQ';
```

```
SEQUENCE_NAME LAST_NUMBER
```

```
-----
```

```
SLNSEQ 1056400
```

You observed this behavior many times and only during peak hours.

What could be the problem and solution for this?

- A. The reserved pool is not configured. Configure the reserved pool.
- B. The keep buffer pool is not configured. Configure the keep buffer pool.
- C. The sequence was created with a CACHE of 100 numbers. Re-create the sequence with a higher CACHE number.
- D. The sequence is aging out of the shared pool. Keep the sequence in the shared pool using the DBMS_SHARED_POOL.KEEP procedure.

Answer: D

Question No : 2

You are working as a DBA on an online transaction processing (OLTP) system. This OLTP system runs on a machine with a single CPU, on which applications primarily perform small random I/Os where each foreground process reads a data block into the buffer cache for updates and the changed blocks are written in batches by the DBWR process. You noticed the increase in I/O requests queued up against a disk, and an increase in the wait time in queue. What would you recommend to improve the latency of I/O requests?

- A. Stripe data across multiple disks.
- B. Increase the size of buffer cache.
- C. Decrease the value of the DBWR_IO_SLAVES parameter.
- D. Increase the value of the DB_WRITER_PROCESSES parameter.

Answer: A

Question No : 3

You received complaints from users about slow response time in some parts of the application. On investigation, you find that this happens whenever the application runs any query against the PROD_DET

table that contains detailed information about all the products.

You gathered the table statistics and executed the following statements to find more information about

this table:

```
SQL> SELECT chain_cnt, 2 round(chain_cnt/num_rows*100,2) pct_chained, 3  
avg_row_len, pct_free ,
```

```
pct_used 4 FROM user_tables 5 WHERE table_name = 'PROD_DET';CHAIN_CNT  
PCT_CHAINED
```

```
AVG_ROW_LEN PCT_FREE PCT_USED-----  
100 3691 10 40 -----4789
```

Which methods can you use to improve the performance? (Choose all that apply.)

- A. online table redefinition
- B. export and import the table
- C. decrease the PCTFREE value
- D. move the table using the ALTER TABLE command

Answer: A,B,D

Question No : 4

What indicates that the rate of writing data blocks to the disk is very slow?

- A. very high cache-hit ratio
- B. increasing number of free buffer waits
- C. increasing number of buffer busy waits
- D. waits on latch: cache buffers chains event

Answer: B

Question No : 5

Your system has been upgraded from Oracle Database 10g to Oracle Database 11g. You imported SQL Tuning Sets (STS) from the previous version that contains important SQL statements. You are not sure about the elapsed time that the SQL statements will have in the new system due to the version change of the optimizer. Which tool provides a comparative report with respect to the elapsed time in both versions of the optimizer?

- A. SQL Access Advisor
- B. SQL Tuning Advisor (STA)
- C. SQL Performance Analyzer (SPA)
- D. Automatic Workload Repository (AWR) compare period report

Answer: C

Question No : 6

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: A,B,D

Question No : 7

Which three statements are true about Statspack usage? (Choose three.)

- A. If both Statspack and AWR are used, they cannot be scheduled to capture snapshots at the same time.
- B. Time Model data cannot be captured by Statspack unless the STATISTICS_LEVEL parameter is set to ALL.
- C. Both Statspack and AWR can be used on the same database only if the Database Diagnostics pack is present.
- D. Statspack can be used to generate a report for a specific instance as well as a specific SQL statement in an instance.
- E. For best performance when running the performance reports, optimizer statistics should be gathered on the Statspack schema.

Answer: C,D,E

Question No : 8

You identified poorly performing SQL by analyzing the V\$SQL and V\$SQLSTATS views. You want to investigate the plans for these SQL statements. Which two methods can you use to pull out the execution plan from the library cache for the already executed SQL? (Choose two.)

- A. Query V\$SQL_PLAN to view the execution plan.

- B. Query DBA_HIST_SQL_PLAN to view the execution plan.
- C. Copy and paste the SQL text from the V\$SQL view and use EXPLAIN PLAN to generate the execution plan.
- D. Use the dbms_xplan.display_cursor function with the SQL ID and child number to generate the execution plan.

Answer: A,D

Question No : 9

View the Exhibit and examine the content.

```
SQL> SELECT tablespace_name, tablespace_size, allocated_space, free_space FROM DBA_TEMP_FREE_SPACE;
```

TABLESPACE_NAME	TABLESPACE_SIZE	ALLOCATED_SPACE	FREE_SPACE
TRMP_L	157286400	144048128	157220864
TEMP	105906176	105906176	0

```
SQL> SELECT (s.tot_used_blocks/f.total_blocks)*100 as pctused
2 FROM (SELECT SUM(used_blocks) tot_used_blocks
3 FROM V$SORT_SEGMENT
4 WHERE tablespace_name='TEMP_L') s,
5 (SELECT SUM(blocks) total_blocks
6 FROM DBA_TEMP_FILES
7 WHERE tablespace_name='TEMP_L') f;
```

PCTUSED
0

ALLOCATED_SPACE and FREE_SPACE indicate that a large amount of space is allocated to the temporary tablespace that is being used. Which action would reclaim the space and refresh these values?

- A. shrinking the tablespace
- B. closing all the current sessions
- C. restarting the database instance
- D. forcing a checkpoint or waiting for a checkpoint to occur

Answer: A

Question No : 10

View the Exhibit and examine the content. ALLOCATED_SPACE and FREE_SPACE are not matching with the TABLESPACE_SIZE value. What would you conclude from this?

Oracle 1z0-054 : Practice Test

```
SQL> SELECT tablespace_name, tablespace_size, allocated_space, free_space FROM D:
```

TABLESPACE_NAME	TABLESPACE_SIZE	ALLOCATED_SPACE	FREE_SPACE
TEMP_L	157286400	144048128	157220864
TEMP	105906176	105906176	0

```
SQL> SELECT (s.tot_used_blocks/f.total_blocks)*100 as pctused
2 FROM (SELECT SUM(used_blocks) tot_used_blocks
3 FROM V$SORT_SEGMENT
4 WHERE tablespace_name='TEMP_L') s,
5 (SELECT SUM(blocks) total_blocks
6 FROM DBA_TEMP_FILES
7 WHERE tablespace_name='TEMP_L') f;
```

```
PCTUSED
-----
0
```

- A. A new temporary file is added to the tablespace.
- B. The temporary space was not sufficient for sort operations.
- C. The temporary space allocated to large sort operations is not deallocated.
- D. TEMP_L is a default temporary tablespace because the numbers show some overhead on this tablespace.

Answer: C

Question No : 11

The users in your online transaction processing (OLTP) environment complain that the response time of the application has increased dramatically. To investigate the problem, view the Exhibit and examine the queries you executed. What do you infer?

Oracle 1z0-054 : Practice Test

```
SQL> SELECT name profile, cnt, decode(total, 0, 0, round(cnt*100/total))
2  percentage
3  FROM (SELECT name, value cnt, (sum(value) over ()) total FROM V$SYSSTAT
5  WHERE name like 'workarea exec%');
```

PROFILE	CNT	PERCENTAGE
workarea executions - optimal	42095	95
workarea executions - onepass	15	0
workarea executions - multipass	1986	5

```
SQL> SELECT * FROM V$PGASTAT;
```

NAME	VALUE	UNIT
aggregate PGA target parameter	20971520	bytes
aggregate PGA auto target	4194304	bytes
global memory bound	154624	bytes
total PGA inuse	108472320	bytes
total PGA allocated	207599616	bytes
maximum PGA allocated	239077376	bytes
total freeable PGA memory	19988480	bytes
process count	55	
max processes count	62	
PGA memory freed back to OS	523763712	bytes
total PGA used for auto workareas	5240832	bytes
maximum PGA used for auto workareas	7198720	bytes
total PGA used for manual workareas	0	bytes
maximum PGA used for manual workareas	537600	bytes
over allocation count	1988	
bytes processed	8550086656	bytes
extra bytes read/written	2.4251E+10	bytes
cache hit percentage	26.06	percent
recompute count (total)	1988	

19 rows selected.

- A. Cursor sharing is not enabled.
- B. The large pool is not configured.
- C. The temporary tablespace is full.
- D. The value set for the SHARED_POOL_SIZE initialization parameter is too small.
- E. The value set for the PGA_AGGREGATE_TARGET initialization parameter is too small.

Answer: E

Question No : 12

Identify the type of recommendation that can automatically be implemented by the Automatic Tuning Advisor as part of automatic SQL tuning task?

- A. recommendation for SQL profiles
- B. recommendation about effective indexing
- C. recommendation about stale or no statistics
- D. recommendation about the structure of SQL statements

Answer: A

Question No : 13

Your company wants to migrate a single instance Oracle 11g Release 1 database to a RAC environment. The database supports middle-tier applications using connection pooling for connecting to database.

Being a part of the performance management team, you want to test the performance of the SQL statements and the peak workload of the applications in the new environment before upgrading the production system.

Given below are some of the steps in random order to accomplish the task:

- 1) Capture the real workload on the existing system.
- 2) Create a SQL Tuning Set (STS) for the statements executed by the application.
- 3) Transport STS on the test environment and use SQL Tuning Advisor to compare the before and after performance.
- 4) Replay the workload on the test environment.
- 5) Analyze the divergence after replay.
- 6) Use Automatic Workload Repository (AWR) to capture the information for top SQL on the existing system.
- 7) Transport STS on the test environment and use SQL Performance Analyzer to compare the before and after performance.

Identify the correct steps required to perform the test.

- A. 2 and 3 only
- B. 6 and 7 only
- C. 1, 4, and 5 only
- D. 2, 3, and 5 only
- E. 2, 3, 4, and 5 only

Answer: C

Question No : 14

View the Exhibit and analyze the output of the query. What does the high value for the table fetch continued rows statistics indicate?

```
SQL> SELECT name,value FROM v$sysstat WHERE name like '%table%';
```

NAME	VALUE
physical reads direct temporary tablespace	502
physical writes direct temporary tablespace	491
DBWR tablespace checkpoint buffers written	18
DBWR transaction table writes	89
transaction tables consistent reads - undo records applied	0
transaction tables consistent read rollbacks	0
auto extends on undo tablespace	0
table scans (short tables)	10782
table scans (long tables)	75
table scans (rowid ranges)	0
table scans (cache partitions)	0
table scans (direct read)	32
table scan rows gotten	10832942
table scan blocks gotten	227752
table fetch by rowid	220813
table fetch continued row	132046
table lookup prefetch client count	0
LOB table id lookup cache misses	0

18 rows selected.

- A. chained rows fetched only
- B. migrated rows fetched only
- C. rows fetched from external tables
- D. both chained and migrated rows fetched
- E. rows fetched from Index Organized Tables (IOT)

Answer: D

Question No : 15

View the Exhibit and examine portions of two ASH reports generated for the same time period. The second report was generated after the ASH data is flushed to disk. Why is the sample count reduced in the second report?