Oracle

Exam 1z0-054

Oracle Database 11g: Performance Tuning

Verson: Demo

[Total Questions: 10]

Question No : 1

Examine the following block that executes a SQL Performance Analyzer task:

```
begin dbms_sqlpa.execute_analysis_task(task_name => 'SPA_SKD4', execution_type =>
'TEST EXECUTE',

execution_name => 'before');

end;
```

Which two statements are true about the execution? (Choose two.)

A. It generates execution plans and execution statistics.
B. It executes only the query part of the DML statements.
C. The test plans are generated and added to plan baseline.
D. The test plans are executed, but not generated if they are available in the plan baseline.

Answer: A,B

Question No : 2

View the Exhibit and examine the sample Automatic Database Diagnostic Monitor (ADDM) report.Which statement is true about the report?

```
Findings and Recommendations
----------------------------

Finding 1: Top SQL by DB Time
Impact is 5.8 active sessions, 100% of total activity.
-------------------------------------------------------
SQL statements consuming significant database time were found.

   Recommendation 1: SQL Tuning
   Estimated benefit is 5.74 active sessions, 98.94% of total activity.
   ---------------------------------------------------------------------
   ....
   ....

   Recommendation 2: SQL Tuning
   Estimated benefit is 5.68 active sessions, 98.04% of total activity.
   ---------------------------------------------------------------------
   Action
      Investigate the SQL statement with SQL_ID "3csh3g3mjhmzh" for possible
      performance improvements.
      Related Object
         SQL statement with SQL_ID 3csh3g3mjhmzh.
         INSERT INTO SPCT VALUES (NULL, 'a')
   Rationale
      SQL statement with SQL_ID "3csh3g3mjhmzh" was executed 1000000 times and
      had an average elapsed time of 0.0024 seconds.
   Rationale
      Waiting for event "buffer busy waits" in wait class "Concurrency"
      accounted for 96% of the database time spent in processing the SQL
      statement with SQL_ID "3csh3g3mjhmzh".


Finding 2: Buffer Busy
Impact is 5.49 active sessions, 94.74% of total activity.
----------------------------------------------------------
Read and write contention on database blocks was consuming significant
database time.

   Recommendation 1: Schema Changes
   Estimated benefit is 5.49 active sessions, 94.74% of total activity.
   ---------------------------------------------------------------------
   ...
   ...|

   Recommendation 2: Schema Changes
   Estimated benefit is 5.49 active sessions, 94.74% of total activity.
   ---------------------------------------------------------------------
   Action
      Consider partitioning the TABLE "SPC.SPCT" with object ID 77025 in a
      manner that will evenly distribute concurrent DML across multiple
      partitions.
      Related Object
         Database object with ID 77025.
   Rationale
      The INSERT statement with SQL_ID "3csh3g3mjhmzh" was significantly
      affected by "buffer busy" waits.
      Related Object
         SQL statement with SQL_ID 3csh3g3mjhmzh.
         INSERT INTO SPCT VALUES (NULL, 'a')
```

A. The sum of impacts under recommendations in Finding 2 cannot exceed 100%.
B. It is possible that the sum of impacts under recommendations in Finding 2 can exceed 100%.
C. The recommendations regarding "buffer busy" wait cannot appear more than once in the report.
D. The sum of impacts of recommendation 2 in Finding 1 and recommendation 2 in Finding 2 cannot exceed 100%.

Answer: B

Question No : 3

You query the V$SYSSTAT view and notice a significantly high value for the redo log space requests statistic.Which three components would you consider for further investigation based on this information? (Choose three.)

A. LGWR
B. checkpoints
C. archiver activity
D. size of the log buffer
E. size of the redo log files

Answer: B,C,E

Question No : 4

Examine the following block:

BEGIN

DBMS_WORKLOAD_REPOSITORY.DROP_BASELINE (baseline_name => 'peak baseline',

cascade => FALSE, dbid => 3310949047);

END;

Why would you use the FALSE value for the CASCADE parameter while dropping the baseline?

A. because it does not remove snapshots associated with the baseline
B. because it does not drop the template that is used to create the baseline
C. because it does not remove the baseline from the performance page if it appears there
D. because it does not remove the Automatic Database Diagnostic Management (ADDM) results from the Automatic Workload Repository (AWR) that were generated using the baseline

Answer: A

Question No : 5

You want to review a summary of all the wait events in an instance since it started over a

month ago.

You issued the following query:

SQL>SELECT event, total_waits, time_waited, average_wait

FROM V$SYSTEM_EVENT

ORDER BY time_waited DESC;

View the Exhibit and examine a portion of the output from the above query.

EVENT	TOTAL WAITS	TIME WAITED	AVERAGE WAIT
rdbms ipc message	759513	6351364	8.36
SQL*Net message from client	241945	2945562	12.17
.....			
.....			
.....			
log file sync	34294	73132	2.13
buffer busy waits	18136	61771	.87
db file sequential read	284405	28963	.05
log file switch (checkpoint incomplete)	1338458	22522	.02
.....			
.....			
.....			

Which events would you consider from the output for further investigation? (Choose all that apply.)

A. log file sync
B. rdbms ipc message
C. SQL*Net message from client
D. log file switch (checkpoint incomplete)

Answer: A,D

Question No : 6

A user, SCOTT, complains that his database session, SID 249, is extremely slow in processing transactions. You check the ADDM report and do not find any overall database performance issues.

You decide to check the wait events, specifically in SCOTT's session and issue the following command:

SQL>SELECT wait_class_id, wait_class, total_waits, time_waited

FROM v$session_wait_class

WHERE sid = 249;

View Exhibit1 and examine the output.

WAIT_CLASS_ID	WAIT_CLASS	TOTAL_WAITS	TIME_WAITED
1893977003	Other	13432	41109
4217450380	Application	18976	381753
3290255840	Configuration	852	2692
3875070507	Concurrency	1498	97963
3386400367	Commit	18356	4561
2723168908	Idle	510426	113736288
2000153315	Network	254544	397
1740759767	User I/O	37309	53822
4108307767	System I/O	101319	9292

You note that the APPLICATION wait class is the second most expensive wait class after the IDLE wait class.

You issue the following query to identify individual waits in the APPLICATION wait class:

SQL>select event, total_waits, time_waited

from v$system_event e, v$event_name n

where n.event_id = e.event_id

and wait_class_id = 4217450380;

View Exhibit2 and examine the output.

EVENT	TOTAL_WAITS	TIME_WAITED
enq: RO - fast object reuse	8	37
enq: TX - row lock contention	2494	388855
SQL*Net break/reset to client	24689	837

Which view would you examine next to pinpoint the problem that is causing this performance issue?

A. DBA_HIST_SESSMETRIC_HISTORY - to find the highest value metrics in the session
B. V$SQLAREA statistics - to find the SQL statements with the highest number of executions
C. V$SESS_TIME_MODEL - to identify whether the SQL statements in the session werespending more of execution time in parsing or waiting for the CPU

D. V$EVENT_HISTOGRAM - to determine whether the TIME_WAITED is evenly distributed across the occurrences of the wait or whether some waits for the events were very long whilst others were short indicating more erratic response times.

Answer: D

Question No : 7

You have a range-partitioned table in your database. Each partition in the table contains the sales data for a quarter. The partition related to the current quarter is modified frequently, whereas other partitions undergo fewer data manipulations. The preferences for the table are set to their default values. You collect statistics for the table using the following command at regular intervals:

SQL> EXECUTE DBMS_STATS.GATHER_TABLE_STATS('SH','SALES',GRANULARITY=>'AUTO');You need statistics to be collected more quickly. How do you achieve this?

A. Set the PUBLISH preference to TRUE for the partition table.
B. Set the NO_VALIDATE preference to TRUE for the partition table.
C. Set the INCREMENTAL preference to TRUE for the partition table.
D. Increase the value of the STALE_PERCENT preference for the partition table.

Answer: C

Question No : 8

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.

```
select * from employees where employee_id=107;

Execution Plan
----------------------------------------------------------
Plan hash value: 1601196873

--------------------------------------------------------------------------------
| Id  | Operation         | Name | Rows  | Bytes | Cost (%CPU)| Time     |
--------------------------------------------------------------------------------
|   0 | SELECT STATEMENT  |      |     1 |    71 |     3   (0)| 00:00:01 |
|*  1 |  TABLE ACCESS FULL| T    |     1 |    71 |     3   (0)| 00:00:01 |
--------------------------------------------------------------------------------

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

   1 - filter("EMPLOYEE_ID"=107)
```

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: C,D

Question No : 9

As a DBA, you notice a change in the performance of your database. Which two pieces of tuning-related information can you examine in the alert log to identify possible causes? (Choose two.)

A. the number of block corruption errors
B. SQL statements that are consuming maximum resources

C. the top five sessions that are consuming maximum resources
D. the instance parameter values that have been changed using ALTER SYSTEM since the last baseline

Answer: A,D

Question No : 10

View the Exhibit.You decide to quickly configure the adaptive threshold. What would happen to the existing metric thresholds shown in the Exhibit?

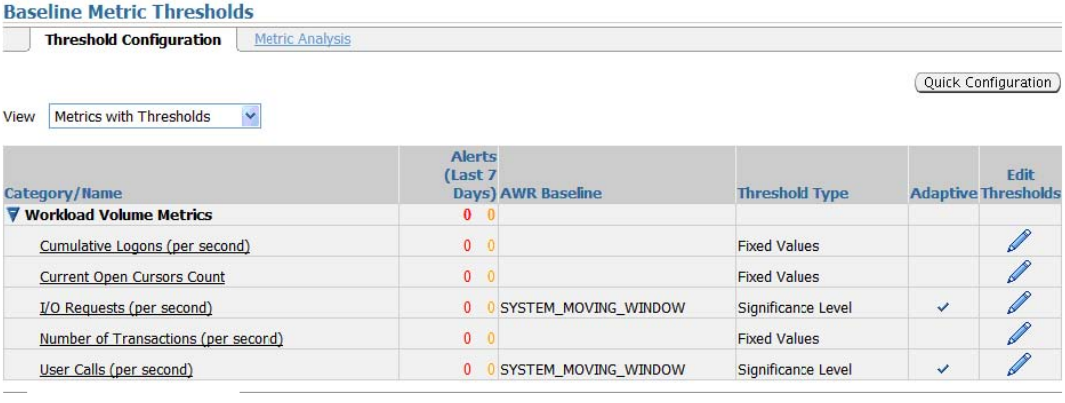

A. It would be overwritten.
B. It would remain in effect.
C. Only the thresholds set with the SYSTEM_MOVING_WINDOW baseline will remain in effect.
D. Only the thresholds set with baselines other than the SYSTEM_MOVING_WINDOW baseline will remain in effect.

Answer: A