Oracle

Exam 1z0-064

Oracle Database 12c: Performance Management and Tuning

Version: 8.0

[Total Questions: 84]

Question No : 1

Examine the parameters set for your database instance:

```
NAME                      TYPE             VALUE
------------------------  --------------   ---------
memory_max_target         big integer      0
memory_target             big integer      0
pga_aggregate_target      big integer      500M
sga_target                big integer      0
db_cache_size             big integer      604M
shared_pool_size          big integer      328M
sga_max_size              big integer      1G
large_pool_size           big integer      24M
```

You upgrade your database to Oracle Database 12*c*. The database supports a mixed workload and works with different workloads at different times. You notice in an ADDM report that the shared pool is inadequately sized. You resize the shared pool by decreasing the sizes of other pools, which results in inadequate sizes for other pools. You want to automate the sizing of SGA components.

Which two actions should you perform? (Choose two.)

A. Set the SGA_TARGET parameter equal to SGA_MAX_SIZE.
B. Set the SGA_TARGET parameter to the sum of DB_CACHE_SIZE,SHARED_POOL, and LARGE_POOL_SIZE.
C. Set the MEMORY_MAX_TARGET parameter to the sum of DB_CACHE_SIZE,SHARED_POOL, and LARGE_POOL_SIZE.
D. Set DB_CACHE_SIZE,SHARED_POOL, and LARGE_POOL_SIZE to their minimum required values.
E. Set the PGA_AGGREGATE_TARGET parameter to 0 and the SGA_TARGET parameter to 1.5G.

Answer: A,E

Question No : 2

Your database supports a mixed workload. In an application, multiple complex queries with functions and expressions are executing. You want to analyze the queries that are currently cached in the library cache to receive recommendations about the usage of indexes and materialized views.

What should you do to achieve this? (Choose the best answer.)

A. Create an STS for the queries cached in the library cache and submit it as an input to SQL Tuning Advisor.
B. Create an STS for the queries cached in the library cache and submit it as an input to SQL Access Advisor.
C. Capture the workload in an STS and submit to SQL Tuning Advisor for recommendations.
D. Create an STS for the queries cached in the library cache and submit it as an input to SQL Performance Analyzer.

Answer: D

Question No : 3

In your database, the measured 99th percentile value is used as the maximum value. You set a warning threshold level of 110% of maximum trigger as an alert.

What is the outcome? (Choose the best answer.)

A. It generates an error because the warning threshold cannot exceed 100%.
B. It generates an error because the percentage of maximum threshold cannot be set with a significance-level threshold value.
C. It generates an alert when an observed metric is 99% of the 99th percentile value as measured over the moving window baseline.
D. It generates an alert when an observed metric is 110% of the 99th percentile value as measured over the moving window baseline.
E. It generates an alert when 1 in 100 observations for an observed metric exceeds the 99th percentile value as measured over the fixed baseline.

Answer: A

Question No : 4

In which three situations can dynamic statistics be used? (Choose three.)

A. when the sampling time is a small fraction of the total time for a query
B. when an execution plan is suboptimal because of complex predicates
C. when extended statistics are not available for SQL statements that use complex

predicates

D. when a query is on a partitioned table with a global index

E. when index statistics are missing on a column that is used in SQL statements with highly selective filters

Answer: A,B,C

Reference: https://docs.oracle.com/database/121/TGSQL/tgsql_statscon.htm#TGSQL341

Question No : 5

Which two statements are true about Compare Period ADDM? (Choose two.)

A. It is automatically invoked whenever the AWR Compare Period report is invoked.

B. It is automatically invoked whenever ADDM is run by default.

C. It verifies if there is any change in the workload or average resource consumption by the SQL executed during the two specified time periods, to ensure 100% accuracy.

D. It can be used to create a comparison report between the Database Replay workload capture report and the replay report.

Answer: C,D

Question No : 6

You are administering a database that supports a DSS workload. Automatic Shared Memory Management is enabled for the database instance. Users issue queries to perform large soft operations and complain about degraded performance of the queries. On investigation, you notice that the queries are performing multipass work area executions and the I/O contention on one of the temporary tablespaces is very high.

Which two can be possible resolutions for this issue? (Choose two.)

A. Increase the size of the large pool.

B. Increase the value of the PGA_AGGREGATE_TARGET parameter.

C. Create a temporary tablespace group and assign it to users.

D. Increase the value of the PGA_AGGREGATE_LIMIT parameter.

E. Create another temporary tablespace and assign it to users.

F. Enable temporary undo.

Answer: C,D

Question No : 7

Which three statements are true about the interpretation of an execution plan? (Choose three.)

A. The cost of the entire plan is indicated by the line with Id 0 and always includes both I/O and CPU resources.
B. The ROWS column indicates the cardinality of each operation and is always calculated by dividing the total number of rows in the table by the number of distinct values in the column used in the WHERE clause predicate.
C. A TABLE ACCESS FULL in the Operation column for a query with a WHERE clause occurs only if no index exists for the filter column.
D. An INDEX UNIQUE SCAN in the Operation column always implies that only one row will be returned.
E. A HASH JOIN in the Operation column always implies that two tables are joined by using an equijoin.

Answer: B,D,E

Question No : 8

You plan to upgrade your production database from Oracle Database 11*g* to 12*c* and also to introduce new objects to the database. You also want to upgrade the hardware. You have already created a test system with the upgrades to be made to the production database. As part of the testing, you want to:

- ✏ analyze and compare the overall database workload with concurrency and transaction characteristics
- ✏ find SQL statements that might get regressed because of the upgrade
- ✏ analyze execution plans for SQL statements for which performance might get regressed
- ✏ analyze the impact of new schema objects on database performance

Which two tools would you recommend to achieve the objective? (Choose two.)

A. Database Replay
B. SQL Tuning Advisor
C. SQL Access Advisor
D. Automatic Database Diagnostic Monitor (ADDM) compare periods report

E. SQL Performance Analyzer
F. Automatic Workload Repository (AWR) compare periods report

Answer: B,E

Question No : 9

Examine the Load Profile section of an AWR report:

	Per Second	Per Transaction	Per Exec	Per Call
DB Time(s):	2.0	0.9	0.02	0.02
DB CPU(s):	0.5	0.2	0.01	0.01
Redo size(bytes):	25,972.2	12,131.8		
Logical reads (blocks):	9,444.6	4,411.6		
Block changes:	144.7	67.6		
Physical reads (blocks):	8,671.9	4,050.7		
Physical writes (blocks):	2,641.5	1,233.9		
User calls:	83.9	39.2		
Parses (SQL):	30.7	14.3		
Hard parses(SQL):	0.4	0.2		
SQL Work Area (MB)	4.6	2.1		
Logons:	2.5	1.2		
Executes (SQL):	88.6	41.4		
Rollbacks:	0.0	0.0		
Transactions:	2.1			

Which two inferences can you derive from the details in this section? (Choose two.)

A. The values for Redo size and Block changes imply that only updates were performed by transactions.
B. The values for Parses (SQL) and Hard parses (SQL) imply that cursor sharing occurred quite often.
C. The values for DB Time and DB CPU imply that the database had a high proportion of idle time during the specified snapshot interval.
D. The values for SQL Work Area and User calls imply that only sort-based operations were performed.
E. The values for Logical reads and Physical reads imply that the number of disk reads per second was less than the total number of DB block reads and consistent gets per second.

Answer: B,D

Question No : 10

You are administering a database that supports a mixed workload. Many applications are running on the middle tier that use connection pools to connect to the database instance. Application users perform OLTP operations during the day and another application performs batch job operations at night. You want to measure and prioritize the two workloads.

Which action would you take to achieve this? (Choose the best answer.)

A. Create database services for the applications, assign individual sessions created by the applications to consumer groups, and then set a priority.
B. Assign profiles to users running the batch operations and make sure that a priority is set for resource limits in profiles.
C. Create database services for the applications and assign different profiles to the sessions to set a relative priority for resource usage.
D. Create database services for the applications, create a job class associated with the service, batch the jobs, and then create jobs by using the job class.

Answer: C

Question No : 11

Your database supports a DSS workload. In an application, a few complex queries that contain multiple functions and expressions are using materialized views. You notice that some queries are performing poorly because they are not benefiting from query rewrites.

Which three actions would you take to improve the performance of queries? (Choose three.)

A. Create an SQL Tuning Set (STS) and submit as input to the SQL Access Advisor to generate recommendations about query rewrite and fast refresh for materialized views.
B. Use the DBMS_MVIEW.EXPLAIN_REWRITE procedure to analyze why a query failed to rewrite.
C. Create an STS and submit as input to the SQL Performance Analyzer to get recommendations about improving the performance of queries.
D. Use the DBMS_ADVISOR.TUNE_MVIEW procedure to get recommendations about rewriting materialized views.
E. Use the DBMS_ADVISOR.QUICK_TUNE procedure to analyze queries based on the usage of query rewrite with materialized views.

Answer: A,C,E

Question No : 12

Examine the partial PLAN_TABLE output:

```
--------------------------------------------------------
Plan hash value: 568005898
--------------------------------------------------------
| Id  | Operation                     | Name     |
--------------------------------------------------------
|    0 | SELECT STATEMENT             |          |
|    1 |  NESTED LOOPS                |          |
|    2 |    TABLE ACCESS BY INDEX ROWID| DEPT    |
|    3 |     INDEX UNIQUE SCAN        | PK_DEPT  |
|    4 |    TABLE ACCESS FULL         | EMP      |
--------------------------------------------------------
```

Which is the correct sequence of execution? (Choose the best answer.)

A. 3, 2, 1, 4, 0
B. 0, 1, 2, 3, 4, 1
C. 0, 4, 1, 3, 2, 1
D. 3, 2, 4, 1, 0
E. 3, 2, 4, 1, 0, 2

Answer: A

Question No : 13

You are administering a database that supports an OLTP workload. An application regularly creates global temporary tables and a large number of transactions are performed on them. You notice that performance is degraded because of excessive generation of undo due to a large number of transactions on the global temporary tables.

What is the recommended action to improve performance? (Choose the best answer.)

A. Increase the size of the undo tablespace and enable undo retention guarantee.
B. Increase the size of the database buffer cache.

C. Enable temporary undo.

D. Increase the size of the temporary tablespace or make it autoextensible.

E. Enable Automatic Segment Space Management (ASSM) for the undo tablespace.

Answer: C

Reference: https://docs.oracle.com/cd/B13789_01/server.101/b10739/undo.htm

Question No : 14

Identify two effects of the DB_FILE_MULTIBLOCK_READ_COUNT parameter on the optimizer. (Choose two.)

A. Decreasing the value of DB_FILE_MULTIBLOCK_READ_COUNT from the default increases the cost of index probes for DSS workloads.

B. A full table scan can become cheaper than index scans if the database instance has a high enough DB_FILE MULTIBLOCK_READ_COUNT for both OLTP and DSS workloads.

C. Increasing the value of DB_FILE_MULTIBLOCK_READ_COUNT within OS limits lowers the costing of an index probe that is done in conjunction with a nested loop for OLTP workloads.

D. In DSS workloads where full table scans may run in parallel and bypass the buffer cache, decreasing the value of DB_FILE_MULTIBLOCK_READ_COUNT from the default increases the cost of full table scans.

E. Increasing the value of DB_FILE_MULTIBLOCK_READ_COUNT within OS limits lowers the cost of full table scans and can result in the optimizer choosing a full table scan over an index scan for both OLTP and DSS workloads.

Answer: B,E

Question No : 15

Which two are prerequisites for running the I/O calibration tool? (Choose two.)

A. The database must be in MOUNT state.

B. The database should be opened in restricted mode.

C. For determining latency time, the STATISTICS_LEVEL parameter must be set to TYPICAL or ALL.

D. The disks to be tested must be configured to use asynchronous I/O for data files.

E. The database instance must be started using an SPFILE.

Answer: C,D

Question No : 16

Your database supports an OLTP system.

Examine the parameter values configured in your database:

sga_max_size = 480M

sga_target = 480M

pga_aggregate_target = 160M

The CUSTOMERS table contains 8,000 rows. The CUST_ID column is the primary key and the COUNTRY_ID column contains only three possible values: 1111, 2222, and 3333.

You execute the commands:

SQL> EXECUTE DBMS_STATS.GATHER_TABLE_STATS('SH','CUSTOMERS');

PL/SQL procedure successfully completed.

SQL> CREATE INDEX COUNTRY_IDX ON CUSTOMERS (COUNTRY_ID);

Index created.

You then perform a series of INSERT, UPDATE, and DELETE operations on the table.

View the Exhibit to examine the query and its execution plan.

```
SQL> SELECT COUNT(*)
 FROM CUSTOMERS
 WHERE COUNTRY_ID = 2222;

 COUNT(*)
----------
      150

SQL> select * from table(dbms_xplan.display_cursor(null,null,'basic rows'));

PLAN_TABLE_OUTPUT
-------------------------------------------------------------------------------
EXPLAINED SQL STATEMENT:
-----------------------
SELECT COUNT(*) FROM CUSTOMERS WHERE COUNTRY_ID = 2222;

Plan hash value: 568322376
--------------------------------------------------------
| ID | Operation            | Name      | Rows   |
--------------------------------------------------------
|    0 | SELECT STATEMENT    |           |        |
|    1 |   SORT AGGREGATE    |           |      1 |
|    2 |    TABLE ACCESS FULL | CUSTOMERS |   8000 |
--------------------------------------------------------
```

Which two options can improve the performance of the query without significantly slowing down the DML operations? (Choose two.)

A. creating a bitmap index on the COUNTRY_ID column
B. regathering statistics on the CUSTOMERS table
C. gathering statistics on the COUNTRY_IDX index
D. creating a histogram on the COUNTRY_ID column
E. increasing the size of the PGA
F. creating an SQL profile
G. creating a KEEP cache

Answer: A,D

Question No : 17

Which four objectives are achieved by using Resource Manager to manage multiple concurrent user sessions that are competing for resources? (Choose four.)

A. distributing available CPU by allocating percentages of CPU time to different users and applications
B. limiting the degree of parallelism of any operation performed by members of a group of users
C. limiting queries based on resource consumption of runaway sessions or calls that consume more than a specified amount of CPU, physical I/O, logical I/O, or elapsed time
D. limiting the number of concurrent sessions for a user

E. limiting the number of user sessions allowed to be concurrently active within a group of users
F. limiting the number of parallel executions that can be executed by a user

Answer: A,B,C,E

Question No : 18

Which three statements are true about using Real-Time Database Operations? (Choose three.)

A. The STATISTICS_LEVEL initialization parameter must be set to ALL to enable automatic SQL monitoring for all long-running queries.
B. The CONTROL_MANAGEMENT_PACK_ACCESS initialization parameter must be set to DIAGNOSTIC+TUNUNG to use Real-Time Database Operations.
C. The STATISTICS_LEVEL initialization parameter can be set to TYPICAL or ALL to enable Real-Time Database Operations.
D. Real-Time Database Operations can be enabled only at the system level.
E. Real-Time Database Operations can be created by using the DBMS_MONITOR or DBMS_SESSION packages.
F. Database operation monitoring starts automatically when a database operation consumes at least five seconds of the CPU or I/O time in a single execution.

Answer: B,C,F

Question No : 19

Which two statements are true about server-generated alerts? (Choose two.)

A. They are always logged in the alert log.
B. They are written to a trace file if the TRACE_ENABLED initialization parameter is set to TRUE.
C. They are generated only when the STATISTICS_LEVEL initialization parameter is set to ALL.
D. They can be generated for user-defined metric thresholds.
E. They appear in the DBA_ALERT_HISTORY view whenever corrective action is taken for an alert.

Answer: D,E

Question No : 20

Examine the parameters set for a database instance supporting a mixed workload:

```
NAME                          TYPE          VALUE
-------------------------     ----------    ------
memory_max_target             big integer   0
memory_target                 big integer   0
pga_aggregate_target          big integer   376M
sga_max_size                  big integer   1G
sga_target                    big integer   0
sort_area_size                integer       65536
```

The database instance supports shared server and dedicated server connections simultaneously. Users complain about increased response times of a few DSS queries. During investigation, you execute the queries:

```
SQL> SELECT d.value as disk, m.value as memory, (d.value/m.value)*100 as ratio
     FROM v$sysstat m, v$sysstat d
     WHERE m.name='sorts (memory)' and d.name='sorts (disk)';
     DISK      MEMORY      RATIO
---------  ---------   --------
     9180     80477    11.40699
SQL> SELECT name,value FROM v$sysstat WHERE name LIKE 'workarea executions%';
NAME                                                      VALUE
------------------------------------------------------   -------
workarea executions - multipass                               89
workarea executions - optimal                              49654
workarea executions - onepass                               1367
```

Based on the output, which two courses of action would you recommend to improve query performance? (Choose two.)

A. Use a parallel hint in the queries.
B. Increase the number of DBW*n* processes.
C. Increase the value of the SORT_AREA_SIZE initialization parameter.
D. Increase the size of the temporary tablespace or add a new temporary tablespace.
E. Increase the value of the PGA_AGGREGATE_TARGET initialization parameter.
F. Increase the size of the large pool.

Answer: C,F

Question No : 21

Examine the parameters set for your database instance:

```
NAME                                 TYPE          VALUE
------------------------------       ------------  -----------
optimizer_capture_sql_plan_baselines boolean       FALSE
optimizer_use_sql_plan_baselines     boolean       TRUE
optimizer_index_cost_adj             integer       100
optimizer_mode                       string        ALL_ROWS
cursor_sharing                       string        EXACT
```

You are administering a database that supports an OLTP workload. Users complain about the degraded performance of some queries. While diagnosing, you notice a large number of hard parses occurring for several syntactically almost identical SQL statements that differ only in literal values in the WHERE clause.

Which two actions would you recommend to improve performance? (Choose two.)

A. Create the KEEP cache and cache the tables used in the queries.
B. Set the CURSOR_SHARING parameter to FORCE.
C. Use bind variables instead of literals.
D. Create SQL plan baselines for the almost identical SQL statements and load them into the cursor cache.
E. Set the OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES parameter to TRUE.

Answer: B,E

Question No : 22

To investigate the slow response time of queries on the TRANS table, you gathered the table statistics and executed the query:

```
SQL> SELECT chain_cnt, round(chain_cnt/num_rows*100,2) pct_chained, avg_row_len,
pct_free , pct_used
     FROM user_tables
     WHERE table_name = 'TRANS';
CHAIN_CNT   PCT_CHAINED AVG_ROW_LEN PCT_FREE  PCT_USED
---------   ----------- ----------- --------  --------
4789              100        3691       10        40
```

The table is stored in a tablespace that has Automatic Segment Space Management (ASSM) enabled. The tablespace is created with a standard block size of 8192 bytes.

Which three can be reasons for the slow response time of the queries? (Choose three.)

A. Row size is too large to fit into a single block during insert operations.
B. Row moves from one data block to another data block because the row grows too large to fit in the original block.
C. The table is subject to frequent insert, update, and delete DML activity leading to sparsely populated blocks.
D. The value of PCTUSED is set to a value lower than the default, causing row changing.
E. The value of PCTFREE is set to a value lower than the default, causing row chaining.

Answer: A,B,C

Question No : 23

You are administering a database that supports multiple applications, which make dedicated connections to the database instance by using different services.

You execute the command to enable tracing of the ORCL1 service:

```
SQL> EXECUTE DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE (service_name => 'ORCL1', WAITS =>
TRUE, BINDS => NULL, instance_name => 'ORCL',plan_stat => NULL);
```

Which two statements are true? (Choose two.)

A. A single trace file is generated for all sessions mapped to the ORCL1 service.
B. SQL trace is enabled for all modules and actions for sessions mapped to the ORCL1 service.
C. An SQL trace file is generated for each session that maps to the ORCL1 service.
D. An SQL trace file is generated for each of the modules using the ORCL1 service.
E. SQL trace is not enabled for the service because a module name is not specified.

Answer: A,C

Question No : 24

For which three problem categories does Automatic Database Diagnostic Monitor (ADDM) provide analysis and recommendations by default? (Choose three.)

A. for network stack-related bandwidth contention
B. for concurrency issues because of buffer busy problems