



Oracle

Exam Questions 1Z0-064

Oracle Database 12c: Performance Management and Tuning

NEW QUESTION 1

You have been asked to use table compression for two large tables. Given are the details of the tables:

The TRANS_DET table:

- ? The table is used by an OLTP application.
- ? High volume insert and update operations are performed on the table.
- ? The table is frequently queried using index range scans.

The TRANS_HISTORY table:

- ? The table is used by a DSS application.
- ? High volume bulk loads are performed on the table.
- ? The table is used to store archival data on which large table full-table scans (FTS) are performed.

Which row store compression would you recommend for these tables with minimal overhead on performance? (Choose the best answer.)

- A. basic table compression for both the tables
- B. advanced row compression for both the tables
- C. basic table compression for the TRANS_HISTORY table and advanced row compression for the TRANS_DET table
- D. basic table compression for the TRANS_DET table and advanced row compression for the TRANS_HISTORY table
- E. warehouse compression for the TRANS_DET table and archive compression for the TRANS_HISTORY table

Answer: A

NEW QUESTION 2

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: BD

NEW QUESTION 3

Examine the partial TOP 10 Foreground Events by Total Wait Time section of an AWR report:

Top 10 Foreground Events by Total Wait Time

Event	Waits	Time (s)	Avg wait (ms)	%Total Call Time	Wait Class
enq: TX - allocate ITL entry	9,799	28,698	2929	32.9	Configurat
db file sequential read	4,827,509	25,964	5	29.7	User I/O
read by other session	2,998,307	18,118	6	20.7	User I/O
CPU time		6,872		7.9	
direct path read	222,425	4,782	21	5.5	User I/O

What should you examine to diagnose the cause of the top three wait events? (Choose the best answer.)

- A. the V\$ACTIVE_SESSION_HISTORY view
- B. the Time Model Statistics section of the AWR report
- C. the SQL statements based on elapsed time from the AWR report
- D. the Latch Activity section
- E. the Segment Statistics section of the AWR report

Answer: B

NEW QUESTION 4

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: CD

NEW QUESTION 5

In the CUSTOMERS table, the values in the CUST_STATE column are dependent on the values in the COUNTRY_ID column. You want to make the optimizer aware of this dependency when these columns are used together in WHERE clause predicates that contain equalities or in-lists. Which two methods achieve this? (Choose two.)

- A. gathering statistics on the CUSTOMERS table and its dependent objects, and then locking the statistics
- B. using SQL plan directives to generate an optimal plan
- C. setting the dynamic statistics level to 4 and setting the OPTIMIZER_USE_PENDING_STATISTICS initialization parameter to true
- D. creating column group statistics, regathering statistics, and ensuring that histograms exist on both these columns

Answer: AD

NEW QUESTION 6

You are administering a database that supports a mixed workload. The CURSOR_SHARING parameter is set to the default value. While analyzing the latest Automatic Workload Repository (AWR) report, you find a large number of cursor: pin S wait on X, cursor: pin X wait on S, and library cache mutex waits in the Top 10 foreground events section. Examine the Instance Efficiency Percentages section in the AWR report:

Instance Efficiency Percentages (Target 100%)

Buffer Nowait %:	100.00	Redo NoWait %:	100.00
Buffer Hit %:	99.95	In-memory Sort %:	100.00
Library Hit %:	62.17	Soft Parse %:	52.72
Execute to Parse %:	47.12	Latch Hit %:	97.95
Parse CPU to Parse Elapsed %:	53.98	% Non-Parse CPU:	70.94

Which three statements are true in this scenario? (Choose three.)

- A. Sessions are waiting for mutexes in share mode on cursors but other sessions are holding the mutexes in exclusive mode.
- B. The CPU is spending more time in finding cursors in the library cache.
- C. Cursors are not getting shared, resulting in a large number of hard parses.
- D. Sessions are waiting for mutexes in exclusive mode on cursors but other sessions are holding the mutexes in share mode.
- E. The buffers required by queries are not found in the buffer cache, thereby increasing expensive disk I/O.

Answer: BDE

NEW QUESTION 7

Users complain about increased response time for queries in your production database that supports an OLTP workload. On investigation, you notice a large number of db file scattered read, latch: cache buffers lru chain, and latch: cache buffers chains wait events: Identify three possible reasons for the increased response time. (Choose three.)

- A. too many sort operations being performed
- B. repeated simultaneous access to a block or small number of blocks
- C. the shared pool is inadequately sized
- D. queries not using indexes and performing full table scans
- E. queries repeatedly fetching blocks that are not in the database buffer cache
- F. cursors are closed explicitly after each execution

Answer: BDE

NEW QUESTION 8

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: ABC

Explanation:

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

NEW QUESTION 9

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
-----
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 DBWn 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: CF

NEW QUESTION 10

You recently joined a new team administering a database.

You notice that full table scans are performing poorly compared with full table scans on the databases you administered in a previous job.

You decide that performance problems are caused by a misconfiguration of factors affecting full table scans.

Which three factors should you investigate to determine the cause of the poorly performing Full Table Scans (FTS)? (Choose three.)

- A. value of DB_FILE_MULTIBLOCK_READ_COUNT
- B. storing query results in the result cache
- C. setting of the DISK_ASYNC_IO parameter to TRUE
- D. setting of the OPTIMIZER_MODE parameter to ALL_ROWS
- E. use of parallel queries
- F. block size of the tablespaces in which the tables being scanned are stored
- G. value of the OPTIMIZER_DYNAMIC_SAMPLING parameter

Answer: ABC

NEW QUESTION 10

Examine the parameters set for your database instance:

NAME	TYPE	VALUE
db_block_size	integer	8192
db_2k_cache_size	big integer	0
db_4k_cache_size	big integer	0
db_8k_cache_size	big integer	0
db_16k_cache_size	big integer	0
db_32k_cache_size	big integer	0

You are asked by a developer to create a table for an application with these requirements:

- ? The table will be used for a DSS application.
 - ? High volume bulk loads will be performed.
 - ? The table will be used to store archival data on which large full-table scans (FTS) will be performed.
- Which attributes are the best for the tablespace in which this table should be created? (Choose the best answer.)

- A. Create it in a locally managed tablespace with ASSM enabled and assign a high value for the PCTFREE attribute.
- B. Create it in a locally managed tablespace with manual segment space management.
- C. Create it in a locally managed tablespace with a bigger nonstandard block size and ASSM enabled.
- D. Create it in locally managed tablespace with ASSM enabled and an additional freelist.

Answer: C

NEW QUESTION 11

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 12c. 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: AE

NEW QUESTION 14

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: DE

NEW QUESTION 19

Which two result in the latest fragmentation within segments and the least external fragmentation within tablespaces? (Choose two.)

- A. automatic segment space managed segments
- B. locally managed tablespaces with uniform extent size
- C. freelist managed segments with one freelist
- D. dictionary managed tablespace
- E. locally managed tablespaces that were converted from dictionary managed tablespaces
- F. freelist managed segments with multiple freelist

Answer: AB

Explanation:

Reference: http://docs.oracle.com/cd/B19306_01/server.102/b14220/logical.htm

NEW QUESTION 24

Your database supports an online transaction processing (OLTP) workload. The database uses ASM storage. One of the ASM disks goes offline because of hardware failure. When the disk is replaced and then added back to the diskgroup, database performance is affected by rebalance operations. Which two actions would you recommend to lower the impact of rebalance operations on the performance of the database? (Choose two.)

- A. Increase the number of ASMB processes.
- B. Decrease the value of the ASM_POWER_LIMIT parameter.
- C. Set the DISK_REPAIR_TIME disk attribute to a lower value.
- D. Specify the POWER clause with a lower value in an ALTER DISKGROUP statement.
- E. Set the DISK_REPAIR_TIME disk attribute to a higher value.

Answer: BD

NEW QUESTION 27

In your database, the locally managed tablespace, USERS, has the default space usage alert set to 85% for the warning level and 97% for the critical level. Which two statements are true? (Choose two.)

- A. Alerts are recorded in both Oracle Enterprise Manager Cloud Control and DBA_OUTSTANDING_ALERTS only when the critical threshold is exceeded.
- B. Alert settings for the warning and critical levels must be disabled before taking the USERS tablespace offline.
- C. Alerts that are triggered are automatically recorded in DBA_ALERT_HISTORY after they are cleared.
- D. Alerts are triggered when the space usage reaches the warning level, again when it reaches the critical level, and yet again when the space usage falls below the critical level.

Answer: BC

NEW QUESTION 32

You are administering a database that supports an OLTP workload. RESULT_CACHE_MODE is set to the default value and a result cache is configured for the instance. Multiple sessions execute syntactically similar queries without dblinks, containing functions and expressions, on tables with no DML activity. Some users complain about poor performance of these queries.

You investigate and find that the queries are frequently performing physical I/O, even though the results fetched by the queries are similar.

Which two actions do you recommend to overcome the problem affecting these queries? (Choose two.)

- A. Set the RESULT_CACHE_MODE parameter to FORCE for the instance.
- B. Use the result cache hint in the queries.
- C. Use bind variables for similar queries instead of literals.
- D. Set the RESULT_CACHE_REMOTE_EXPIRATION parameter to a nonzero value.
- E. Configure the KEEP pool and cache the queried tables used in the KEEP pool.

Answer: AB

NEW QUESTION 35

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: ABCE

NEW QUESTION 38

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: CD

NEW QUESTION 41

Examine the partial Activity Over Time section of an Active Session History (ASH) report:

Slot Time (Duration)	Slot Count	Event	Event Count	% Event
14:10:50 (1.2 min)	5	control file sequential read	4	0.11
		CPU + Wait for CPU	1	0.03
14:12:00 (3.0 min)	9	CPU + Wait for CPU	5	0.14
		control file parallel write	2	0.05
		null event	1	0.03

Which two inferences are correct? (Choose two.)

- A. In the first time slot, five different sampled sessions were connected to the database instance.
- B. In the second time slot, out of the nine sampled sessions connected to the database instance, only one sampled session was idle at the time of report generation.
- C. In the first time slot, only one sampled session was using the CPU.
- D. In the second time slot, five different sampled sessions were using the CPU.
- E. In the second time slot, 0.14% of the time was spent on the CPU.

Answer: AE

NEW QUESTION 43

A senior DBA asks you to decrease the values of the connect_time_scale and think_time_scale replay processing parameters to 50 to preprocess the workload for replay.

What three could be reasons for this change? (Choose three.)

- A. to reduce the elapsed time between two successive user calls from a session.
- B. to decrease the number of concurrent users during replay
- C. to increase the number of concurrent users during replay
- D. to reduce the time of replay
- E. to decrease the wait for a query, caused by noncommitted transactions

Answer: CDE

NEW QUESTION 44

Examine the command to change a parameter value from the default to 50: SQL> ALTER SYSTEM SET OPTIMIZER_INDEX_COST_ADJ = 50;

What is the effect of changing the value of the parameter? (Choose the best answer.)

- A. It influences the optimizer to use full table scans instead of index scans as the estimated cost of full table scan is reduced.
- B. It influences the optimizer to use bitmap indexes as the estimated cost of conversion from bitmap to rowid is reduced.
- C. It influences the optimizer to always use fast full index scans as the estimated cost of using an index is reduced.
- D. It influences the optimizer to use indexes instead of full table scans as the estimated cost of using an index is reduced.

Answer: A

Explanation:

Reference: http://www.dba-oracle.com/oracle_tips_cost_adj.htm

NEW QUESTION 45

In which three situations does DB time always increase? (Choose three.)

- A. when the host is CPU bound for foreground processes
- B. when I/O wait time increases for foreground processes
- C. when more connections are made to a database instance
- D. when CPU consumption by background processes increases
- E. when wait time for data to be sent over a network increases

Answer: ABC

Explanation:

Reference: <http://www.oracle.com/technetwork/oem/db-mgmt/s317294-db-perf-tuning-with-db-time-181631.pdf> (page 21)

NEW QUESTION 50

Your database supports multiple applications. The applications run on the middle tier and use connection pooling for connecting to the database. You notice that the sessions created by the applications are competing for resources. You want to statistically measure the workload and set priorities. What action must you perform to achieve this? (Choose the best answer.)

- A. Create services for the applications and set a relative priority by assigning them to application users and using the DBMS_MONITOR.SERV_MOD_ACT_TRACE_ENABLE procedure to trace the services.
- B. Create services for the applications and set a relative priority by assigning them to application users and using the DBMS_MONITOR.SESSION_TRACE_ENABLE procedure to trace the services.
- C. Create services for the applications and set the relative priority of services within an instance by mapping the services directly to consumer groups.
- D. Create services for the applications and set a relative priority by assigning them to application users.

Answer: A

NEW QUESTION 55

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

NEW QUESTION 60

Examine the query and its output:

```
SQL> SELECT sid, seq#, event, p1text, p1, p2text, p2, p3text, p3, wait_time,
seconds_in_wait, state FROM v$session_wait WHERE sid = 24;
```

SID	SEQ#	EVENT	P1TEXT	P1	P2TEXT	P2	P3TEXT	P3	WAIT_TIME
24	104	db file scattered read	file#	12	block#	1221	blocks	8	-1

Which two inferences can be definitely derived from this output? (Choose two.)

- A. The db file scattered read event has occurred 104 times in this session for file# 12.
- B. The session has completed performing a full table scan.
- C. The SQL statements in this session are performing excessive disk reads.
- D. The multiblock factor is 8 for this I/O but it could vary for the other I/O events.

Answer: AC

NEW QUESTION 65

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

NEW QUESTION 69

Examine the parameters set for a database instance:

NAME	TYPE	VALUE
memory_max_target	big integer	0
memory_target	big integer	0
lock_sga	boolean	FALSE
pre_page_sga	boolean	TRUE
sga_max_size	big integer	1G
sga_target	big integer	1G
result_cache_max_size	big integer	0
result_cache_mode	string	MANUAL

An application performs a large number of identical queries on small lookup tables very frequently. Users complain about the slow response time of queries on these tables. On investigation, you notice that buffers are getting aged out of the buffer cache. To mitigate the issue, you increase the value of the SGA_MAX_SIZE and SGA_TARGET parameters, but after some time, you notice the same issue again. Which two would you recommend as long-term solutions for this issue? (Choose two.)

- A. increasing the size of the database buffer cache
- B. configuring Automatic Memory Management
- C. configuring the KEEP buffer pool and altering tables to use the KEEP pool
- D. pinning the cursors of the queries in the library cache
- E. configuring the result cache for the instance

Answer: AB

NEW QUESTION 72

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

- A. It analyzes the performance of a database instance based on the time period covered by the most recent AWR snapshot, and generates recommendations based on hard-coded criteria.
- B. It can analyze performance issues that occurred in past events provided they fall within the AWR retention period.
- C. ADDM resource utilization and cost of analysis depends on the actual load on the database and the number of performance problems analyzed.
- D. It first identifies the performance symptoms, and then refines them to reach the root cause with the singular aim of reducing the DB CPU metric.
- E. It documents only those components and wait classes that are significantly impacting the performance of the database.

Answer: AB

NEW QUESTION 77

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
- C. for high-load PL/SQL execution and compilation, and high-load Java usage
- D. for application-level lock contention.

Answer: BCD

NEW QUESTION 81

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: AD

NEW QUESTION 85

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 several INSERT, UPDATE, and DELETE operations, significantly altering the data in 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 three options would improve the performance of the query? (Choose three.)

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

Answer: ABF

NEW QUESTION 90

You are administering a database that supports an OLTP workload. CURSOR_SHARING is set to EXACT for the instance. An application is frequently executing almost identical queries that vary in literal values in the WHERE clause, causing a large number of hard parses to occur.

Which four statements would be true if you use bind variables for these queries? (Choose four.)

- A. Mutex contention in the library cache will be reduced.
- B. The optimizer will use one parent cursor and one child cursor for each SQL statement with different literal values.
- C. Hard parses will be reduced for the queries.
- D. The optimizer will use bind peeking and subsequent execution of the queries will always generate the same plans irrespective of the cardinality.
- E. The optimizer will generate the same plan for all bind values if no histograms exist on the columns used in the WHERE clause of these queries.
- F. The optimizer will use bind peeking and use the literal value to determine the execution plan for these queries.

Answer: ACDE

NEW QUESTION 94

Which two statements are true about ADDM or Real-Time ADDM? (Choose two.)

- A. ADDM can be run manually by selecting any range of AWR snapshots available within the AWR retention period, provided they do not cover a time period when

the instances were restarted.
 B. ADDM runs in Partial mode to analyze any hung database issues.
 C. Real-Time ADDM can proactively detect and diagnose transient performance issues that last for a few seconds.
 D. Real-Time ADDM is automatically invoked by ADDM at the end of every hour.

Answer: AC

NEW QUESTION 97

Examine the parameters set for a database instance:

NAME	TYPE	VALUE
memory_max_target	big integer	0
memory_target	big integer	0
pga_aggregate_target	big integer	256M
sga_max_size	big integer	1G
sga_target	big integer	1G

The database supports a mixed workload. Users complain about the increased response time of a few DSS queries. During investigation, you execute the query:

```
SQL> SELECT name,value FROM v$sysstat WHERE name LIKE 'workarea executions%';
NAME                                VALUE
-----
workarea executions - multipass      557
workarea executions - optimal       47256
workarea executions - onepass       1146
```

Based on the output, which two are possible ways to improve the performance of the queries? (Choose two.)

- A. Enable temporary undo.
- B. Enable Automatic Memory Management.
- C. Increase the number of DBWn processes.
- D. Enable Automatic Shared Memory Management.
- E. Increase the value of the SGA_TARGET parameter.
- F. Increase the value of the PGA_AGGREGATE_TARGET parameter.

Answer: CE

NEW QUESTION 101

You are administering a database that supports an OLTP workload. An application performs a large number of small transactions. Users complain about increased response times for transactions. On investigation, you find that the cache hit ratio is 69%. Examine a partial output from V\$SYSTEM_EVENT:

EVENT	TOTAL_WAITS
buffer busy waits	103500
read by other session	795497
free buffer waits	76398

Which four can be possible reasons for the increased response time? (Choose four.)

- A. The database buffer cache is inadequately sized.
- B. DBWR is not writing the dirty buffers fast enough.
- C. A large number of blocks are fetched from disks frequently.
- D. Several full table scans are performed by transactions.
- E. Blocks are aging out of the buffer cache frequently.
- F. Many sessions are waiting for buffers that are currently being read into the buffer cache by other sessions.

Answer: ABEF

NEW QUESTION 103

.....

Thank You for Trying Our Product

We offer two products:

1st - We have Practice Tests Software with Actual Exam Questions

2nd - Questions and Answers in PDF Format

1Z0-064 Practice Exam Features:

- * 1Z0-064 Questions and Answers Updated Frequently
- * 1Z0-064 Practice Questions Verified by Expert Senior Certified Staff
- * 1Z0-064 Most Realistic Questions that Guarantee you a Pass on Your First Try
- * 1Z0-064 Practice Test Questions in Multiple Choice Formats and Updates for 1 Year

100% Actual & Verified — Instant Download, Please Click
[Order The 1Z0-064 Practice Test Here](#)