

1Z0-064 Dumps

Oracle Database 12c: Performance Management and Tuning

<https://www.certleader.com/1Z0-064-dumps.html>



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

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: ACE

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 statements are true about viewing the details of Real-Time Database Operations? (Choose two.)

- A. In V\$SQL_MONITOR monitoring, statistics are cumulative over several executions of the SQL statement that is being monitored in a session.
- B. SQL Developer can be used to view running database operations.
- C. Oracle Enterprise Manager Database Express can be used to view running database operations.
- D. When the SQL statement that is being monitored is executing, V\$SQL_MONITOR is refreshed once every minute.
- E. After the execution ends, the monitoring information in V\$SQL_MONITOR is deleted immediately.
- F. Oracle Enterprise Manager Cloud Control can be used to view running database operations.

Answer: AD

NEW QUESTION 5

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: ABC

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

You plan to upgrade your production database from Oracle Database 11g to 12c 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: BE

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:

NAME	TYPE	VALUE
parallel_degree_policy	string	MANUAL
workarea_size_policy	string	AUTO
sort_area_size	integer	65536
memory_max_target	big integer	0
memory_target	big integer	0
pga_aggregate_target	big integer	256M
sga_target	big integer	1G

Your database supports a mixed workload and users have dedicated server connections. Users complain about the increased response time of a few queries that are performing large sort operations. On investigation, you notice an increase in the number of multipass work area executions and high number of direct path write wait events.

Which two actions could improve the performance? (Choose two.)

- A. increasing the value of the SORT_AREA_SIZE parameter
- B. increasing the value of the PGA_AGGREGATE_TARGET parameter
- C. enabling Automatic Memory Management for the instance
- D. increasing the size of the default temporary tablespace
- E. using parallel hint in queries performing large sort operations
- F. enabling Automatic Shared Memory Management for the instance

Answer: AF

NEW QUESTION 10

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 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 14

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 16

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 17

Which three methods can you use to create a pre-change SQL trial to capture performance data by using SQL Performance Analyzer? (Choose three.)

- A. executing SQL statements in an SQL Tuning Set (STS) on a test database by using database links to the production database.

- B. generating only execution plans on a test database without actually running SQL statements.
- C. generating an execution plan and statistics for selective SQL statements captured in an STS
- D. loading performance data and execution plans from an STS.
- E. generating both execution plans and statistics for each SQL statement in an STS by actually running the SQL statements on a test database.

Answer: BDE

Explanation: Reference: https://docs.oracle.com/cd/E11882_01/server.112/e41481/spa_pre_change.htm#RATUG1_81

NEW QUESTION 21

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 26

You define the warning threshold for the tablespace usage metric for the USERS tablespace to be 60% and the critical threshold to be 80%. Which two sources should you check for the alert information when either the warning or the critical threshold is exceeded? (Choose two.)

- A. the alert log
- B. Oracle Enterprise Manager Cloud Control
- C. DBA_ALERT_HISTORY
- D. DBA_OUTSTANDING_ALERTS
- E. DBA_ACTIVE_SESSION_HISTORY
- F. DBA_THRESHOLDS

Answer: AF

NEW QUESTION 31

Which two statements are true about DB time in V\$SYS_TIME_MODEL? (Choose two.)

- A. DB time cannot exceed the total elapsed time (wall clock time) since the database instance started.
- B. DB time cannot exceed the maximum number of concurrent sessions multiplied by the actual elapsed time for each session.
- C. DB time includes the time spent on client processes and background processes.
- D. Reducing DB time allows a database instance to support more user requests by using the same resources.
- E. DB time is always greater than or equal to the DB CPU time.

Answer: DE

NEW QUESTION 35

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 40

Which two statements are true about the interpretation of Buffer Cache Hit Ratio in the Instance Efficiency Percentages section of an AWR report? (Choose two.)

- A. A high value indicates that the buffer cache is adequately sized for the current workload.
- B. Poor hit ratios indicate that a large number of indexed lookups or small table scans are being performed.
- C. A low hit ratio does not necessarily imply that increasing the size of the buffer cache will improve performance.
- D. A high hit ratio may indicate that repeated scanning of the same large table or index is being performed.
- E. A low hit ratio indicates that a KEEP buffer pool should be configured based on the size of the largest object accessed in the buffer cache.

Answer: CD

NEW QUESTION 44

You are administering a database that supports an OLTP workload. The CURSOR_SHARING parameter is set to EXACT for the instance. The performance of queries issued by one of the modules has degraded. The queries executed by the module are almost identical in syntax. To investigate, you analyze the latest AWR report and find a large number of latch:shared pool wait events and also a high percentage of the hard parse elapsed time. Which two can be reasons for this? (Choose two.)

- A. The I/O performance is slow.
- B. Bind variables are not used for similar queries, causing hard parses.
- C. Repeated access to a small number of blocks.
- D. Excessive time is spent on finding cached cursors in the library cache.
- E. The CURSOR_SHARING parameter is set to EXACT, which does not allow similar queries to share a cursor.

Answer: BC

NEW QUESTION 46

Examine the parameters set for your database instance:

NAME	TYPE	VALUE
optimizer_capture_sql_plan_baselines	boolean	TRUE
optimizer_use_sql_plan_baselines	boolean	TRUE

You notice that for one particular SQL statement, the optimizer generates a new better plan than the plans in the SQL Plan Management Base. Which action is taken by the optimizer? (Choose the best answer.)

- A. It adds the newly generated plan as an accepted but non-fixed plan.
- B. It adds the newly generated plan as enabled and accepted.
- C. It adds the newly generated plan as enabled but not accepted.
- D. It adds the newly generated plan as a fixed plan, which will be used each time the SQL statement is executed.

Answer: B

NEW QUESTION 50

Examine the Time Model Statistics section of an AWR report:

Statistic Name	Time (s)	% of DB Time
sql execute elapsed time	12,416.14	86.45
DB CPU	9,223.70	64.22
parse time elapsed	935.61	6.51
hard parse elapsed time	884.73	6.16
failed parse elapsed time	21.39	.72
PL/SQL execution elapsed time	153.51	1.07
hard parse (sharing criteria) elapsed time	25.96	0.18
connection management call elapsed time	14.00	0.10
hard parse (bind mismatch) elapsed time	4.74	0.03
PL/SQL compilation elapsed time	1.20	0.01
repeated bind elapsed time	0.22	0.00
sequence load elapsed time	0.11	0.00
DB time	14,362.96	
background elapsed time	731.00	
background cpu time	72.00	

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

- A. The available CPU resources were not utilized to their maximum capacity.
- B. All sequence numbers used during this AWR time interval were cached.
- C. A large number of connected user sessions were idle.
- D. New child cursors were created because of new bind values or usage of literal values as well as different bind types or sizes.
- E. The DB CPU time was not spent exclusively for processing SQL statements.

Answer: DE

NEW QUESTION 54

You observe that queries are performing poorly on the SALES_RECORDS table in your database. On investigation, you find that at the end of each day the contents of the SALES_RECORDS table are moved to the SALES_HISTORY table. The delete operations cause the table to be sparsely populated. The SALES_RECORDS table is created in a tablespace using Automatic Segment Space Management (ASSM) and row movement is enabled. The table must be accessible 24x7.

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

- A. Perform EXPORT, DROP, and IMPORT operations on the SALES_RECORDS table.
- B. Shrink the SALES_RECORDS table by using the ALTER TABLE...SHRINK SPACE command.
- C. Move the SALES_RECORDS table to a different location by using the ALTER TABLE...MOVE command.
- D. Deallocate the space in the SALES_RECORDS table by using the ALTER TABLE...DEALLOCATE UNUSED command.
- E. Move the SALES_RECORDS table to a tablespace by using manual segment space management.
- F. Reorganize the SALES_RECORDS table online by using the DBMS_REDEFINITION package.

Answer: BD

NEW QUESTION 57

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 61

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 62

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: BE

NEW QUESTION 63

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 64

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 66

Examine this list of possible tasks:

1. Ensure that STATISTICS_LEVEL is set to TYPICAL or ALL.
2. Ensure that TIMED_STATISTICS is set to TRUE.
3. Set MAX_DUMP_FILE_SIZE to UNLIMITED and DIAGNOSTIC_DEST to an appropriate destination.
4. Ensure that SQL_TRACE is set to TRUE.
5. Enable tracing at the database instance level by using the DBMS_MONITOR.DATABASE_TRACE_ENABLE procedure.

6. Enable tracing in the required session by using the DBMS_SESSION.SET_SQL_TRACE procedure.
7. Run TKPROF with the EXPLAIN parameter on the output trace file.
8. Run the trcsess utility on the output trace files, and then run TKPROF on the output of the trcsess utility.
Select the minimum tasks to perform, in the correct order, to generate both a formatted trace file with timing information and an explain plan for each SQL statement for all sessions. (Choose the best answer.)

- A. 1, 2, 5, 8
- B. 1, 3, 6, 7
- C. 2, 4, 5, 8
- D. 1, 3, 4, 5, 6, 7
- E. 1, 2, 4, 8

Answer: C

NEW QUESTION 67

Examine the partial AWR report taken for a time period of 60 minutes:

Top 10 Foreground Events by Total Wait Time

Event	Waits	Time (s)	Avg wait (ms)	%Total Call Time	Wait Class
resmgr: cpu quantum	475,956	152,859	320	75.2	Scheduler
CPU time		47,880		23.5	
db file sequential read	3,374,890	16,868	5	7.8	User I/O
db file scattered read	196,265	4,278	22	2.1	User I/O
log file sync	177,735	4,579	29	5.4	Commit
.....					
.....					
.....					

Operating System Statistics DB/Inst: ****/**** Snaps: 56708/56709
Statistic Total

.....
BUSY_TIME 5,707,832
IDLE_TIME 2

.....
NUM_CPUS 32

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

- A. The database user calls are issuing frequent explicit commits.
- B. The CPUs are busy executing server processes and background processes for a considerable amount of CPU time.
- C. The database user calls are spending most of their time in I/O for single block reads.
- D. The database user calls are spending most of their time waiting for sessions that are in more important consumer groups.

Answer: BC

NEW QUESTION 71

Examine the partial TKPROF output for an SQL statement:

```
SQL> SELECT city_id
      FROM city_names
      WHERE code = 'DLR'?
```

call	count	cpu	elapsed	disk	query	current	rows
Parse	1	0.06	0.10	0	0	0	0
Execute	1	0.02	0.02	0	0	0	0
Fetch	1	0.23	0.30	31	31	3	1

```
Misses in library cache during parse: 0
Parsing user id: 02 (USER2)
```

Rows	Execution Plan
0	SELECT STATEMENT
2340	TABLE ACCESS (BY ROWID) OF 'CITY_NAMES'
0	INDEX (RANGE SCAN) OF 'CITY_NAMES_NAME' (NON-UNIQUE)

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

- A. Array fetch operations were not performed for this query.
- B. No hard parse was performed for this query.
- C. The number of logical I/Os is almost equal to the number of physical I/Os.
- D. Another transaction held a shared lock on the table, thereby causing a significant delay.

Answer: BD

NEW QUESTION 76

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 80

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 81

You want to capture AWR data to monitor performance variation every Monday between 9:00 AM and 12:00 PM for three months and automatically remove the older AWR data every fortnight. How would you achieve this? (Choose the best answer.)

- A. Create AWR baselines.
- B. Create SQL plan baselines.
- C. Create repeating baseline templates.
- D. Create database services and make sure that user connections use them to connect to the database instance.
- E. Create a single baseline template.

Answer: D

NEW QUESTION 86

You are administering a database that supports a mixed workload. You upgrade your database from Oracle Database 11g to 12c and after the upgrade, users

complain about degraded performance of some queries. The SQL plan baselines imported from the previous version are present for the queries and are loaded to the SQL Management Base as accepted plans. On further investigation, you find that better plans are generated but not used by the optimizer. Examine the parameters set for the instance:

NAME	TYPE	VALUE
optimizer_capture_sql_plan_baselines	boolean	FALSE
optimizer_use_sql_plan_baselines	boolean	TRUE

Which three tasks would you perform to improve the performance of these queries? (Choose three.)

- A. Gather statistics for the objects used in the queries.
- B. Use the DBMS_SPM.EVOLVE_SQL_PLAN_BASELINE function to evolve new plans and fix the plans for the statements.
- C. Create an SQL Tuning Set (STS) and run it through the SQL Access Advisor to generate recommendations.
- D. Create an STS and run it through the SQL Tuning Advisor to generate recommendations.
- E. Set the OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES parameter to TRUE.
- F. Use the DBMS_SPM.ALTER_SQL_PLAN_BASELINE function to alter the accepted plans as fixed plans.

Answer: ABC

NEW QUESTION 87

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: CD

NEW QUESTION 90

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 93

Examine the initialization parameters set for a database instance:

NAME	TYPE	VALUE
dbwr_io_slaves	integer	0
db_writer_processes	integer	1
filesystemio_options	string	NONE
disk_asynch_io	boolean	TRUE

The database supports an OLTP workload. Applications connect to the instance using shared server connections and perform small, random I/Os. All the data files are on the same disk. You notice free buffer wait events for sessions in the database instance. To solve the problem, you increase the size of the buffer cache. But after some time, you notice sessions waiting again on free buffer waits. What will you recommend to alleviate the issue? (Choose the best answer.)

- A. Run the I/O calibration tool.
- B. Configure the database instance to make asynchronous I/O available to DBWR.
- C. Spread the data files over multiple disks, controllers, and I/O buses to ensure that there are no hotspots in the I/O subsystem.
- D. Configure dedicated server connections for the applications.

Answer: B

NEW QUESTION 97

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 99

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 100

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 101

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 105

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

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

NEW QUESTION 106

.....

Thank You for Trying Our Product

* 100% Pass or Money Back

All our products come with a 90-day Money Back Guarantee.

* One year free update

You can enjoy free update one year. 24x7 online support.

* Trusted by Millions

We currently serve more than 30,000,000 customers.

* Shop Securely

All transactions are protected by VeriSign!

100% Pass Your 1Z0-064 Exam with Our Prep Materials Via below:

<https://www.certleader.com/1Z0-064-dumps.html>