1Z0-064 Dumps

Oracle Database 12c: Performance Management and Tuning

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

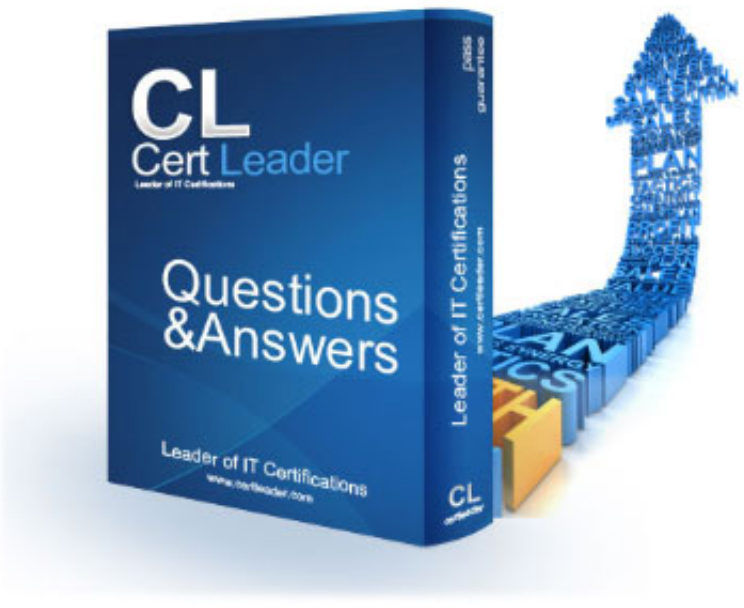

NEW QUESTION 1
Examine the Load Profile section of an AWR report:

	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 2
You are administering a database that supports an OLTP workload. Users complain about the degraded response time of a query. You want to gather new statistics for objects accessed by the query and test query performance with the new statistics without affecting other sessions connected to the instance.
The STALE_PERCENT statistic preference is set to a default value and the STATISTICS_LEVEL parameter is set to TYPICAL.
Which two actions would you take to accomplish the task? (Choose two.)

A. Set the STALE_PERCENT statistic preference to a higher value than the default, and then gather statistics.
B. Set the STATISTICS_LEVEL parameter to ALL for the instance.
C. Set the INCREMENTAL preference to TRUE, and then gather statistics.
D. Set the OPTIMIZER_USE_PENDING_STATISTICS parameter to TRUE for the session in which you want to test the query.
E. Set the PUBLISH statistic preference to FALSE, and then gather statistics.
F. Set the NO_INVALIDATE statistic preference to TRUE, and then gather statistics.

Answer: BE

NEW QUESTION 3
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 4
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

NEW QUESTION 5
You plan to upgrade your production database from Oracle Database 11g to 12c. As part of the upgrade, you want to introduce new indexes and materialized views. You have already created a test system with Oracle Database 12c, having the same structure and data as the production database, along with new schema objects to be added to the production database.
You want to identify regressed SQL statements, if any, which may have been caused by schema changes and the change in the optimizer version.
Which two methods would you use to achieve this? (Choose two.)

A. Create an SQL Tuning Set (STS) for the SQL statements on the production database and submit as input to the SQL Tuning Advisor on the test database.
B. Create an STS for the SQL statements on the production database and submit as input to the SQL Performance Analyzer with the OPTIMIZER_FEATURES_ENABLE parameter first set to 11.2.0.1, and then to 12.1.0.1 on the test database.
C. Generate an Automatic Workload Repository (AWR) compare periods report with snapshots taken before and after schema changes on the test database.
D. Capture the production database workload, replay it on the test system by using Database Replay, and analyze by using the workload replay compare period report.
E. Create an STS for the SQL statements on the production database and submit as input to the SQL Access Advisor on the test database.
F. Create an STS for the SQL statements on the production database before and after changes and submit as input to the SQL Performance Analyzer on the test database.

Answer: AD

NEW QUESTION 6
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 7
Your database supports an OLTP workload during the day and batch processing at night. You want to monitor performance metrics to detect any degradation of performance in both types of workloads over a time period of 30 days.
Examine this list of possible steps:
1. Create a fixed baseline.
2. Create a baseline template.
3. Create a new moving window baseline.
4. Increase the retention period default value to 30 days.
5. Increase the size of the existing moving window baseline to 30 days.
6. Create warning and critical alerts for the relevant metrics.
7. Enable adaptive thresholds to detect the workload patterns and specify a high- significance-level threshold type.
8. Enable adaptive thresholds to detect the workload patterns and set different threshold values as a percentage of the maximum value.
Which option represents the required steps in the correct order? (Choose the best answer.)

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

Answer: E

NEW QUESTION 8
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 9
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 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 structure of the EMPLOYEES table.

```
SQL> desc employees

Name                                      Null?    Type
----------------------------------------- -------- ----------------
EMPLOYEE_ID                               NOT NULL NUMBER(6)
FIRST_NAME                                         VARCHAR2(20)
LAST_NAME                                 NOT NULL VARCHAR2(25)
HIRE_DATE                                 NOT NULL DATE
JOB_ID                                    NOT NULL VARCHAR2(10)
SALARY                                             NUMBER(8,2)
MANAGER_ID                                         NUMBER(6)
DEPARTMENT_ID                                      NUMBER(4)
```

EMPLOYEE_ID is the primary key. No other indexes exist on this table. View the Exhibit to examine the commands and their output.

```
SQL> select department_id, count(department_id) from employees group by
department_id order by 2;

DEPARTMENT_ID  COUNT(DEPARTMENT_ID)
-------------  --------------------
           40                     1
           10                     1
           70                     1
           20                     2
          110                     2
           90                     3
           60                     5
           30                     6
          100                     6
           80                    34
           50                    45
11 rows selected.

SQL> var dept_id number
SQL> exec :dept_id := 50
SQL> select count(*) from employees where department_id= :dept_id;
   COUNT(*)
----------
        45
SQL> /
   COUNT(*)
----------
        45

SQL> SELECT CHILD_NUMBER, IS_BIND_SENSITIVE AS "BIND_SENSI", IS_BIND_AWARE AS
"BIND_AWARE", IS_SHAREABLE AS "BIND_SHARE" FROM V$SQL
WHERE SQL_TEXT LIKE 'select count(*) from emp%';

CHILD_NUMBER   BIND_SENSI   BIND_AWARE   BIND_SHARE
------------   ----------   ----------   ----------
           0   N            N            Y
```

Which two actions should you perform to make the cursor bind aware? (Choose two.)

A. Create a histogram on the DEPARTMENT_ID column.
B. Change the default CURSOR_SHARING value to FORCE.
C. Execute the query with the same DEPARTMENT_ID value multiple times.
D. Create an index on the DEPARTMENT_ID column.
E. Gather statistics for the index.
F. Regather statistics on the table.

Answer: CD

NEW QUESTION 15
For your database some users complain about not being able to execute transactions. Upon investigation, you find that the problem is caused by some users performing long- running transactions that consume huge amounts of space in the UNDO tablespace.
You want to control the usage of the UNDO tablespace only for these user sessions. How would you avoid the issue from repeating in future? (Choose the best answer.)

A. Create a profile for the users with the LOGICAL_READS_PER_SESSION and LOGICAL_READS_PER_CALL limits defined.
B. Create external roles to restrict the usage of the UNDO tablespace and assign them to the users.
C. Set the threshold for UNDO tablespace usage for the users.
D. Implement a Database Resource Manager plan by mapping the users to a resource consumer group with limits defined for UNDO tablespace usage.

Answer: D

NEW QUESTION 20
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 25
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 30
Examine the output of the query executed to diagnose the reason for performance degradation of queries:

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

NAME                                                             VALUE
----------------------------------------------------------       ----------
physical reads direct temporary tablespace                          50
physical writes direct temporary tablespace                        491
DBWR tablespace checkpoint buffers written                          18
DBWR transaction table writes                                       89
transaction tables consistent reads - undo records applied           0
transaction tables consistent read rollbacks                        0
auto extends on undo tablespace                                     0
table scans (short tables)                                      10782
table scans (long tables)                                          75
table scans (rowid ranges)                                          0
table scans (cache partitions)                                      0
table scans (direct read)                                          32
table scan rows gotten                                        10832942
table scan blocks gotten                                       4227752
table fetch by rowid                                           2220813
table fetch continued row                                      1132046
table lookup prefetch client count                                  0
LOB table id lookup cache misses                                    0
```

Which three factors will you investigate further to identify the cause of the performance degradation? (Choose three.)

A. Check the number of disk sorts.
B. Check for the causes of the full table scans.
C. Check the number of chained or migrated rows.

Answer: ABC

NEW QUESTION 32
You execute this query twice in a session:

```
SQL>select product_name
from order_items o, product_information p
where o.unit_price = 15 and quantity > 1
and p.product_id = o.product_id;
```

Then you query V$SQL_SHARED_CURSOR for details about child cursors as shown.

```
SQL>select c.child_number, c.use_feedback_stats from v$sql_shared_cursor c
where c.sql_id = 'an4zdfz0h7513';

CHILD_NUMBER   USE_FEEDBACK_STATS
------------   ------------------
           0   Y
           1   N
```

Which two statements are true? (Choose two.)

A. No statistics were collected during the first execution of the query.
B. A subsequent execution of the query in this session is likely to undergo a soft parse.
C. The second execution of the query was hard parsed because the estimated cardinality was inaccurate.
D. A subsequent execution of the query in this session will undergo a hard parse.
E. The second execution of the query was hard parsed because extended statistics were collected after the first execution of the query.

Answer: BC

NEW QUESTION 35
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 40
Which two statements are true about Active Session History (ASH)? (Choose two.)

A. The Data Sample size available in an ASH report is dynamic and, at any given moment, is directly related to the amount of work being performed.
B. ASH contains sampled data from all sessions that are connected to a database instance at any given moment.
C. ASH samples data from V$SESSION every second.
D. An ASH report can be used to identify the service that may be the cause of a transient performance problem.

Answer: AD

NEW QUESTION 41
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 42
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 44
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 49
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 50
Your database supports a mixed workload. The ERP application creates short sessions and performs small, random I/Os; the REPORTING application executes long-running DSS queries.
You want to set a priority for the workload generated by the ERP application and optimize resource usage for them.
Which three objectives can be achieved by the Resource Manager? (Choose three.)

A. limiting the amount of time that a session is idle and blocking other sessions of the ERP application
B. limiting the amount of undo generated by operations performed by sessions created by the ERP application
C. creating two resource plans with resource limits defined for the workload generated by the applications and automatically changing resource plans based on the workload
D. allocating a lower percentage of CPU to sessions used by the REPORTING application than to those used by the ERP application
E. limiting the physical I/O performed by the sessions or users of the ERP application that are connected to the database

Answer: BDE

NEW QUESTION 54
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 58
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 62
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 63
Examine the query and its output:

```
SQL>select sid,state,wait_time/100 "WAIT TIME IN SECONDS", event from v$session where
username='HR';
```

Output:

```
SID         STATE               WAIT TIME IN SECONDS     EVENT
-------     -----------         --------------------     -------------
2832        WAITED KNOWN TIME      2029                  rdbms ipc message
3346        WAITING                   0                  enq: TX - row lock
contention
4208        WAITING                   0                  SQL*Net message from client
```

Which two statements are true? (Choose two.)

A. Session 2832 had to wait 2029 seconds for a message to arrive because of a network bottleneck.
B. Session 4208 is either idle or experiencing poor response time due to a network or resource bottleneck on the client process.
C. Session 3346 is in wait state because it wants to lock a row in a block in which other sessions have already locked rows, and there is no free ITL slot available in this block.
D. Session 3346 is in wait state because either it is waiting to update a row that is locked by another session or another session is trying to insert the same key value in a UNIQUE index.
E. Session 4208 is definitely idle and should be killed to free network resources.

Answer: AD

NEW QUESTION 67
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 68
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 73
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 74
For which two requirements can you always use the V$ACTIVE_SESSION_HISTORY view? (Choose two.)

A. to investigate intermittent performance problems in a session, only when the problem lasted less than five minutes in the last twelve hours
B. to find the exact number of executions of a specific query in a session in the last two minutes
C. to identify which module in an application was being executed in a session
D. to identify a scheduler job that is not responding
E. to find the amount of Program Global Architecture (PGA) memory being currently consumed by a session

Answer: CE

NEW QUESTION 75
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 79
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 80
Which two actions should you take to monitor the throughput generated by the modules of an application? (Choose two.)

A. Use the Resource Manager.
B. Enable SQL Trace at the session level.
C. Create a service.
D. Use a dedicated server configuration.
E. Use the DBMS_APPLICATION_INFO package to define the current module and action so that they appear in V$SESSION.

Answer: BE

NEW QUESTION 81
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 85
Which two actions can reduce library cache latch contention for an OLTP application that repeatedly executes queries containing a mix of literals and bind variables? (Choose two.)

A. setting the OPEN_CURSORS parameter to hold a sufficient number of concurrently open cursors
B. coding the application such that a frequently executed statement is parsed only onceand executed repeatedly as required
C. setting the CURSOR_SHARING parameter to EXACT
D. avoiding the granting of privileges on objects associated with cursors during peak load
E. enabling Automatic Memory Management and allocating at least 50% of the available memory for SHARED_POOL_SIZE
F. configuring shared server connections

Answer: BE

Explanation: Reference: http://docs.oracle.com/cd/B28359_01/server.111/b28274/memory.htm

NEW QUESTION 87
Which two situations can lead to sparsely populated index blocks? (Choose two.)

A. Data is frequently inserted using direct path load into a table with an index.
B. Indexed columns in a table are frequently updated.
C. Values in an indexed column are inserted using monotonically incrementing sequences.
D. Bulk delete operations are performed on a table with indexes.
E. Online table move operations are performed frequently on a table with indexes.

Answer: BD

NEW QUESTION 89
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 90
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 95

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 100
.....

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