



Oracle

Exam 1z0-054

Oracle Database 11g: Performance Tuning

Version: 5.0

[Total Questions: 192]

Question No : 1

You notice that two wait events, Latch:cache buffer chains and Latch:cache buffer LRU chains, appear consistently in the Top 5 Timed Events for your database. Which two options describe what these wait events indicate? (Choose two.)

- A. Latch:cache buffer LRU chains - indicates excessive block replacement
- B. Latch:cache buffer LRU chains - indicates insufficient size of application cursor cache
- C. Latch:cache buffer chains - indicates inefficient SQL that accesses many blocks repeatedly
- D. Latch:cache buffer chains - indicates that there are several sessions waiting for a buffer that is presently being read into the buffer cache by another session

Answer: A,C

Question No : 2

You are working on an online transaction processing (OLTP) system. By day most of the application users perform queries accessing the most recently added or modified rows. The applications have most of the queries based on multiple tables. But at night, some batch processing is also done. Which two actions would you recommend to choose a goal for the optimizer based on the needs of your application? (Choose two.)

- A. setting the OPTIMIZER_MODE parameter to ALL_ROWS at the instance level
- B. setting the OPTIMIZER_MODE parameter to FIRST_ROWS_n at the instance level
- C. asking the developer to add a hint /*ALL_ROWS*/ in the long-running batch processing queries
- D. asking the developer to add a hint /*FIRST_ROWS_n*/ in the long-running batch processing queries

Answer: B,C

Question No : 3

You are a DBA in ABC Corp. You are working on a DSS system.. The applications directly connect to the database to perform transactions. The application users also perform direct loads. The company wants to upgrade from Oracle Database 10g version 10.2.0.1 to

Oracle Database 11g. You want to test the performance of the SQL statements in the new environment before upgrading the production database. To accomplish this, you set up the test database with Oracle Database 11g, with the same data and schemas as in production.

What would you recommend to accomplish the task?

- A. gathering all the SQL statements and using SQL Tuning Advisor on the test database for performance analysis
- B. capturing the workload on the production database, replaying it on the test machine, and comparing the performance
- C. creating a SQL Tuning Set by capturing the SQL workload on the production database and using SQL Performance Analyzer for performance analysis
- D. configuring AWR to capture the maximum number of SQL statements, transporting it to the test system, and using SQL Access Advisor for performance analysis

Answer: C

Question No : 4

View Exhibit1 and examine the indexes on the CUSTOMERS table.

```
SQL> SELECT i.table_name, column_name, o.object_name , o.object_id, i.index_type
FROM   user_objects o,
       user_indexes i,
       user_ind_columns c
WHERE  o.object_type= 'INDEX'
AND    i.table_name LIKE '%CUST%'
AND    i.index_name = o.object_name
AND    i.index_name = c.index_name
```

TABLE_NAME	COLUMN_NAME	OBJECT_NAME	OBJECT_ID	INDEX_TYPE
CUSTOMERS	CUST_GENDER	CUSTOMERS_GENDER_BIX	70685	BITMAP
CUSTOMERS	CUST_MARITAL_STATUS	CUSTOMERS_MARITAL_BIX	70686	BITMAP
CUSTOMERS	CUST_ID	CUSTOMERS_PK	70473	NORMAL
CUSTOMERS	CUST_YEAR_OF_BIRTH	CUSTOMERS_YOB_BIX	70687	BITMAP
CUSTOMERS	COUNTRY_ID	CUST_COUNTRY	79047	NORMAL

The statistics for the CUSTOMERS table have been updated recently by using the following command:

```
SQL> EXEC
DBMS_STATS.GATHER_TABLE_STATS('SH','CUSTOMERS',method_opt=>'FOR ALL
INDEXED COLUMNS SIZE AUTO');
```

View Exhibit2 to examine a query plan. Even though the index is present on the

COUNTRY_ID and CUST_GENDER columns, the query uses a full table scan. What could be the reason?

```
SQL> SELECT i.table_name, column_name, o.object_name , o.object_id, i.index_type
FROM   user_objects o,
       user_indexes i,
       user_ind_columns c
WHERE  o.object_type= 'INDEX'
AND    i.table_name LIKE '%CUST%'
AND    i.index_name = o.object_name
AND    i.index_name = c.index_name
```

TABLE_NAME	COLUMN_NAME	OBJECT_NAME	OBJECT_ID	INDEX_TYPE
CUSTOMERS	CUST_GENDER	CUSTOMERS_GENDER_BIX	70685	BITMAP
CUSTOMERS	CUST_MARITAL_STATUS	CUSTOMERS_MARITAL_BIX	70686	BITMAP
CUSTOMERS	CUST_ID	CUSTOMERS_PK	70473	NORMAL
CUSTOMERS	CUST_YEAR_OF_BIRTH	CUSTOMERS_YOB_BIX	70687	BITMAP
CUSTOMERS	COUNTRY_ID	CUST_COUNTRY	79047	NORMAL

- A. because the histogram statistics for the COUNTRY_ID column are not updated
- B. because the DB_FILE_MULTIBLOCK_READ_COUNT initialization parameter is set to a high value
- C. because the optimizer calculates the cost of accessing blocks by using a full table scan to be less as compared to index scans, even though indexes are available
- D. because indexes on CUST_GENDER and COUNTRY_ID columns are of different types, the index on the CUST_GENDER column is bitmap index, and on COUNTRY_ID columns is btree index.

Answer: C

Question No : 5

You are working on an online transaction processing (OLTP) system. You are investigating the reasons for performance degradation in the database.

View the Exhibit named TOP5EVENTS and note the top five events.

Top 5 Timed Foreground Events

Event	Waits	Time(s)	Avg wait (ms)	% DB time	Wait Class
DB CPU		500		52.98	
free buffer waits	14,517	158	11	16.73	Configuration
db file sequential read	4,349,775	146	0	15.46	User I/O
log file sync	1,462	19	13	2.00	Commit
enq: KO - fast object checkpoint	54	14	266	1.52	Application

View the Exhibit named TOPSQL and note the problems related to these statements.

SQL ordered by Reads

- Total Disk Reads: 4,452,178
- Captured SQL account for 99.5% of Tot.

Physical Reads	Executions	Reads per Exec	%Total	CPU Time (s)	Elapsed Time (s)	SQL Id	SQL Module	SQL Text
4,228,364	11	384,396.73	94.97	191.30	269.66	3amqadvdn93cc	SQL*Plus	DECLARE CURSOR C2 IS SELECT...
4,198,589	109,989	38.17	94.30	162.28	248.83	431mwkyt65jbb	SQL*Plus	SELECT SUM(AMOUNT_SOLD) TOTAL ...
66,019	52	1,269.60	1.48	152.04	229.28	8mfzykm4g68th	SQL*Plus	DECLARE CURSOR C2 IS SELECT...
51,159	112	456.78	1.15	11.28	50.76	9a258qtr5d5mh	SQL*Plus	DECLARE max_records NU...
49,255	12,359	3.99	1.11	18.28	40.64	658qfzar410kx	SQL*Plus	SELECT ORDER_ID FROM (SELECT O...
47,521	111	428.12	1.07	50.11	154.35	1fvwvlysp4psd	SQL*Plus	DECLARE max_orders NUM...
29,268	113	259.01	0.66	18.26	41.18	4ju491r28v70d	SQL*Plus	DECLARE new_order_id NUMBER(1...
26,962	52	518.50	0.61	2.83	19.99	7z52r0tb5wxmm	SQL*Plus	SELECT * FROM ORDER_ITEMS
22,229	792	28.07	0.50	9.80	14.38	0u9dmvxy318w0	SQL*Plus	SELECT SUM(AMOUNT_SOLD) TOTAL ...
21,960	2,752	7.98	0.49	4.08	29.10	21v0ddi14q7mk	SQL*Plus	DELETE FROM ORDER_ITEMS WHERE ...

View the Exhibit named INSTACT and note the table scans.

Instance Activity Statistics

- [Instance Activity Stats](#)
- [Instance Activity Stats - Absolute Values](#)
- [Instance Activity Stats - Thread Activity](#)

[Back to Top](#)

Instance Activity Stats

Statistic	Total	per Second	per Trans
Batched IO (bound) vector count	19	0.03	0.00
Batched IO (full) vector count	19	0.03	0.00
Batched IO block miss count	3,957	7.05	0.09
table scans (direct read)	64	0.11	0.00
table scans (long tables)	124	0.22	0.00
table scans (short tables)	56,541	100.76	1.32

Which is the most appropriate solution?

- Create indexes on short tables.
- Use the CACHE hint to access long tables.
- Keep the long tables in the keep buffer pool.
- Keep the short tables in the keep buffer pool.
- Keep the short tables in the recycle buffer pool.

Answer: D

Question No : 6

View the Exhibit and examine the content. ALLOCATED_SPACE and FREE_SPACE are not matching with the TABLESPACE_SIZE value. What would you conclude from this?

```
SQL> SELECT tablespace_name, tablespace_size, allocated_space, free_space FROM D:
```

TABLESPACE_NAME	TABLESPACE_SIZE	ALLOCATED_SPACE	FREE_SPACE
TEMP_L	157286400	144048128	157220864
TEMP	105906176	105906176	0

```
SQL> SELECT (s.tot_used_blocks/f.total_blocks)*100 as pctused
2 FROM (SELECT SUM(used_blocks) tot_used_blocks
3 FROM V$SORT_SEGMENT
4 WHERE tablespace_name='TEMP_L') s,
5 (SELECT SUM(blocks) total_blocks
6 FROM DBA_TEMP_FILES
7 WHERE tablespace_name='TEMP_L') f;
```

```
PCTUSED
-----
0
```

- A. A new temporary file is added to the tablespace.
- B. The temporary space was not sufficient for sort operations.
- C. The temporary space allocated to large sort operations is not deallocated.
- D. TEMP_L is a default temporary tablespace because the numbers show some overhead on this tablespace.

Answer: C

Question No : 7

Examine the SQL statement given below:

```
SQL> EXPLAIN PLAN for
```

```
SELECT department_name,count(*)
```

```
FROM hr.employees e, hr.departments d
```

```
WHERE e.department_id=d.department_id
```

```
GROUP BY d.department_name;
```

To view the execution plan for the statement, you queried plan_table. View the Exhibit and examine the execution plan. Which two statements are true about the execution plan?

(Choose two.)

```
SQL> SELECT id "id",parent_id "par", position "pos",
        lpad(' ',2*level)||operation||decode(id,0,'cost='||POSITION) "operation",
        options "option", object_name "object", object_node "table queue",
        other_tag "parallel oper type", distribution "row dist",other "slave SQL"
FROM plan_table
CONNECT by prior id=parent_id START WITH id=0
ORDER BY id
```

id	par	pos	operation	option	object	table	quen	parllel	oper	type	row	dist	slave SQL
0		4	SELECT STATEMENT	cost=4									
1	0	1	HASH		GROUP BY								
2	1	1	NESTED LOOPS										
3	2	1	TABLE ACCESS	FULL	DEPARTMENTS								
4	2	2	INDEX	RANGE SCAN	EMP_DEPARTME								NT_IX

- A. The execution plan is displayed for the executed SQL statement.
- B. The EXPLAIN PLAN command does not actually execute the statement.
- C. The row with the id column having a value of 4 is the first step executed in the execution plan.
- D. The row with the id column having a value of 0 is the first step executed in the execution plan.

Answer: B,C

Question No : 8

In which two scenarios would you recommend using RAID level 5 as an option for database storage?(Choose two.)

- A. For a database with lower recovery (MTTR) overhead
- B. For a database in which the applications primarily perform read operations and predominantly sequential I/O is performed
- C. For a database in which performance is the primary goal of the database because applications are demanding in terms of I/O operations
- D. For a database in which the database size grows rapidly, the cost of storage is critical and data accessibility is more important than performance

Answer: B,D

Question No : 9

View the Exhibit exhibit1 to examine the series of SQL commands and parameter settings.

```
SQL> SHOW PARAMETER OPTIMIZER
```

NAME	TYPE	VALUE
optimizer_capture_sql_plan_baselines	boolean	TRUE
optimizer_dynamic_sampling	integer	2
optimizer_features_enable	string	11.1.0.6
optimizer_index_caching	integer	0
optimizer_index_cost_adj	integer	100
optimizer_mode	string	ALL_ROWS
optimizer_secure_view_merging	boolean	TRUE
optimizer_use_invisible_indexes	boolean	FALSE
optimizer_use_pending_statistics	boolean	FALSE
optimizer_use_sql_plan_baselines	boolean	TRUE

```
SQL> SELECT * FROM sh.sales WHERE quantity_sold > 40 ORDER BY prod_id;
SQL> SELECT * FROM sh.sales WHERE quantity_sold > 40 ORDER BY prod_id;
SQL> ALTER SESSION SET OPTIMIZER_MODE=FIRST_ROWS;
SQL> SELECT * FROM sh.sales WHERE quantity_sold > 40 ORDER BY prod_id;
```

View the Exhibit exhibit2 to examine the plans available in the SQL plan baseline.

Select Name	SQL Text	Enabled	Accepted	Fixed	Auto Purge	Created	Last Modified
<input type="checkbox"/> SYS_SQL_PLAN_89447021cf314e9e	select * from hr.employees where job_id='CLERK'	YES	YES	NO	YES	Jul 20, 2008 7:02:30 PM	Jul 20, 2008 7:16:48 PM
<input type="checkbox"/> SYS_SQL_PLAN_891170210572d2c8	select * from hr.employees where job_id='CLERK'	YES	NO	NO	YES	Jul 20, 2008 7:20:45 PM	Jul 20, 2008 7:20:45 PM
<input type="checkbox"/> SYS_SQL_PLAN_7ed8568135b3cdca	SELECT NAME NAME COL PLUS SHOW PARAM,DECODE (TYPE,1...	YES	YES	NO	YES	Jul 21, 2008 2:40:44 PM	Jul 21, 2008 2:40:44 PM
<input type="checkbox"/> SYS_SQL_PLAN_4698b35ddf463620	select * from table(dbms_xplan.display (null,null,'...	YES	YES	NO	YES	Jul 20, 2008 7:04:22 PM	Jul 20, 2008 7:04:22 PM
<input type="checkbox"/> SYS_SQL_PLAN_467a776254bc8843	select * from sh.sales where quantity_sold > 40 or...	YES	YES	NO	YES	Jul 21, 2008 2:25:42 PM	Jul 21, 2008 2:25:42 PM
<input type="checkbox"/> SYS_SQL_PLAN_467a776211df68d0	select * from sh.sales where quantity_sold > 40 or...	YES	NO	YES	YES	Jul 21, 2008 2:41:22 PM	Jul 21, 2008 2:41:56 PM

The first plan (in red) is created when OPTIMIZER_MODE is set to ALL_ROWS and the second plan (in blue) is created when OPTIMIZER_MODE is set to FIRST_ROWS. Which SQL plan baseline would be used if the SQL query in exhibit1 is executed again when the value of OPTIMIZER_MODE is set to FIRST_ROWS?

- A. the second plan, because it is a fixed plan
- B. the first plan, because it is an accepted plan
- C. the second plan, because it is the latest generated plan in FIRST_ROW mode
- D. A new plan, because the second plan in FIRST_ROW mode is not an accepted plan

Answer: B

Question No : 10

As part of a proactive tuning activity, you have the Automatic Workload Repository (AWR) and ADDM enabled for your database with the default setting of snapshot retention as 8 days and snapshot interval as 60 minutes. Which performance issue does NOT appear in the ADDM and needs alternative tools for investigation?

- A. suboptimal parameter settings
- B. application-level block contention
- C. performance issues related to the Resource Manager
- D. short-lived or intermittent performance issues lasting for less than 5 minutes

Answer: D

Question No : 11

View the Exhibit and analyze the output of a query executed on the V\$MEMORY_RESIZE_OPS view. Why do INITIAL_SIZE, TARGET_SIZE, and FINAL_SIZE columns have the value zero for some of the components?

```
SQL> SELECT component, oper_type, oper_mode, initial_size, target_size, final_size, status
2 FROM v$memory_resize_ops;
```

COMPONENT	OPER_TYPE	OPER_MODE	INITIAL_SIZE	TARGET_SIZE	FINAL_SIZE	STATUS
shared pool	STATIC		0	142606336	142606336	COMPLETE
large pool	STATIC		0	4194304	4194304	COMPLETE
ASM Buffer Cache	STATIC		0	0	0	COMPLETE
java pool	STATIC		0	12582912	12582912	COMPLETE
streams pool	STATIC		0	0	0	COMPLETE
SGA Target	STATIC		0	192937984	192937984	COMPLETE
DEFAULT buffer cache	STATIC		0	16777216	16777216	COMPLETE
KEEP buffer cache	STATIC		0	8388608	8388608	COMPLETE
RECYCLE buffer cache	STATIC		0	0	0	COMPLETE
DEFAULT 2K buffer cache	STATIC		0	0	0	COMPLETE
DEFAULT 4K buffer cache	STATIC		0	0	0	COMPLETE
DEFAULT 8K buffer cache	STATIC		0	0	0	COMPLETE
DEFAULT 16K buffer cache	STATIC		0	0	0	COMPLETE
DEFAULT 32K buffer cache	STATIC		0	0	0	COMPLETE
Shared IO Pool	STATIC		0	0	0	COMPLETE
PGA Target	STATIC		0	192937984	192937984	COMPLETE
java pool	SHRINK	DEFERRED	12582912	8388608	8388608	COMPLETE
DEFAULT buffer cache	GROW	DEFERRED	16777216	20971520	20971520	COMPLETE

18 rows selected.

- A. because they are the only components that are auto-tuned
- B. because they are the only components that are not auto-tuned
- C. because they are the only components that have not undergone any manual resizing
- D. because they are the only components that do not have a default size or their respective parameters are not set in the initialization parameter file

Answer: D

Question No : 12

You work as a consultant DBA for various clients. A performance issue in one of the online transaction processing (OLTP) systems is reported to you and you received the Automatic Workload Repository (AWR) report generated in the database. The main sections of the AWR report is shown in the Exhibits.

View the Exhibit named DBTIME and note % DB Time.

Top 5 Timed Foreground Events

Event	Waits	Time(s)	Avg wait (ms)	% DB time	Wait Class
DB CPU		108		94.00	
db file sequential read	315,309	5	0	4.62	User I/O
latch: shared pool	1,227	0	0	0.41	Concurrency
db file parallel read	186	0	1	0.16	User I/O
control file sequential read	3,935	0	0	0.05	System I/O

View the Exhibit named TIMEMODEL and note what has contributed to % DB Time.

Time Model Statistics

- Total time in database user-calls (DB Time): 114.8s
- Statistics including the word "background" measure background process time, and so do not contribute to the DB time statistic
- Ordered by % of DB time desc. Statistic name

Statistic Name	Time (s)	% of DB Time
sql execute elapsed time	111.41	97.02
DB CPU	107.95	94.00
parse time elapsed	85.03	74.05
hard parse elapsed time	73.55	64.05
PL/SQL execution elapsed time	5.83	5.08
PL/SQL compilation elapsed time	2.20	1.92
hard parse (sharing criteria) elapsed time	0.46	0.40
connection management call elapsed time	0.44	0.38
hard parse (bind mismatch) elapsed time	0.27	0.24
sequence load elapsed time	0.06	0.05
repeated bind elapsed time	0.02	0.01
DB time	114.83	
background elapsed time	2.90	
background cpu time	0.95	

View the Exhibit named EFFICIENCY and examine the various percentages shown.

Instance Efficiency Percentages (Target 100%)

Buffer Nowait %:	100.00	Redo NoWait %:	100.00
Buffer Hit %:	74.48	In-memory Sort %:	100.00
Library Hit %:	68.42	Soft Parse %:	34.91
Execute to Parse %:	3.01	Latch Hit %:	99.95
Parse CPU to Parse Elapsed %:	0.01	% Non-Parse CPU:	34.94

Which option describes the correct understanding of the main problem?

- A.** The Exhibit DBTIME shows that the DB CPU event consumed very high % DB Time, which indicates that the CPU is very slow.
- B.** sql execute elapsed time and parse time elapsed are consuming very high % DB Time in the TIMEMODEL Exhibit. This indicates too many soft parses.
- C.** hard parse elapsed time has the majority time in parse time elapsed, which is shown in the TIMEMODEL Exhibit. This indicates inadequate database buffer cache.
- D.** % Non-Parse CPU: and Soft Parse %: are low in the EFFICIENCY Exhibit, which shows that very few statements are found in the cache. This indicates that there were too many hard parses.

Answer: D

Question No : 13

You are working on an online transaction processing (OLTP) system. The middle-tier applications use connection pooling to connect to the database. Presently, you have a single-node database. The company plans to migrate the database to a RAC environment. Before you move to a RAC environment, you want to test the performance of the SQL statements and the peak workload on the new environment. To accomplish the Database Capture for replay, you identified the peak workload period on the existing system and started the Database Capture. Which client requests to the database can be captured as part of the workload capture? (Choose all that apply.)

- A.** flashback query
- B.** distributed transactions
- C.** logging in and logging out of sessions
- D.** all DDL statements having bind variables
- E.** direct path load of data from external files

Answer: C,D

Question No : 14

View the Exhibit and examine the content of the V\$MEMORY_RESIZE_OPS view. Which two conclusions can you draw? (Choose two.)

COMPONENT	OPER_TYPE	OPER_MODE	INITIAL_SIZE	TARGET_SIZE	FINAL_SIZE	STATUS
ASH Buffer Cache	STATIC		0	0	0	COMPLETE
PGA Target	STATIC		0	188743680	188743680	COMPLETE
Shared IO Pool	STATIC		0	0	0	COMPLETE
DEFAULT 32K buffer cache	STATIC		0	0	0	COMPLETE
DEFAULT 16K buffer cache	STATIC		0	0	0	COMPLETE
DEFAULT 8K buffer cache	STATIC		0	0	0	COMPLETE
shared pool	STATIC		0	146800640	146800640	COMPLETE
large pool	STATIC		0	4194304	4194304	COMPLETE
java pool	STATIC		0	12582912	12582912	COMPLETE
streams pool	STATIC		0	0	0	COMPLETE
SGA Target	STATIC		0	188743680	188743680	COMPLETE
DEFAULT buffer cache	STATIC		0	8388608	8388608	COMPLETE
KEEP buffer cache	STATIC		0	8388608	8388608	COMPLETE
RECYCLE buffer cache	STATIC		0	0	0	COMPLETE
DEFAULT 2K buffer cache	STATIC		0	0	0	COMPLETE
DEFAULT 4K buffer cache	STATIC		0	0	0	COMPLETE
shared pool	GROW	IMMEDIATE	146800640	150994944	150994944	PENDING
SGA Target	GROW	IMMEDIATE	188743680	192937984	192937984	COMPLETE
PGA Target	SHRINK	IMMEDIATE	188743680	184549376	184549376	COMPLETE
DEFAULT buffer cache	GROW	DEFERRED	8388608	12582912	12582912	COMPLETE
shared pool	SHRINK	DEFERRED	150994944	146800640	146800640	COMPLETE
DEFAULT buffer cache	SHRINK	MANUAL	12582912	8388608	8388608	COMPLETE
large pool	GROW	MANUAL	4194304	8388608	8388608	COMPLETE

- A. Automatic Memory Management (AMM) is enabled.
- B. Automatic Shared Memory Management (ASMM) is enabled but not AMM.
- C. The Shared Pool grew in IMMEDIATE mode to avoid an out-of-memory error.
- D. The Default Buffer Cache grew in DEFERRED mode to avoid an out-of-memory error.
- E. The memory is transferred in IMMEDIATE mode only if there are entirely empty granules available in other components.

Answer: A,C

Question No : 15

You work as a DBA for a company and as a performance improvement measure, you implemented the result cache in your database. Many users in the company say that performance has improved on the queries they use, but some users complain that they have not got any performance benefit on the queries they use.

You checked all the queries they use and the following is one of them:

```
SQL> SELECT /*+ RESULT_CACHE */ slnoq.currval as "SLNO", prod_id, pdname, 2  
cust_name FROM sales WHERE sl_date < sysdate;
```

View the Exhibit and examine the testing performed to check this.

```
SQL> EXECUTE DBMS_RESULT_CACHE.FLUSH;
```

```
PL/SQL procedure successfully completed.
```

```
SQL> EXECUTE DBMS_RESULT_CACHE.MEMORY_REPORT;
```

```
Result Cache Memory Report
```

```
[Parameters]
```

```
Block Size = 1K bytes
```

```
Maximum Cache Size = 1376K bytes (1376 blocks)
```

```
Maximum Result Size = 68K bytes (68 blocks)
```

```
[Memory]
```

```
Total Memory = 5132 bytes [0.004% of the Shared Pool]
```

```
... Fixed Memory = 5132 bytes [0.004% of the Shared Pool]
```

```
... Dynamic Memory = 0 bytes [0.000% of the Shared Pool]
```

```
PL/SQL procedure successfully completed.
```

```
SQL> SELECT /*+ RESULT_CACHE */ slnoq.currval as "SLNO", prod_id, pdname, cust_name
```

```
2 FROM sales
```

```
3 WHERE sl_date < sysdate;
```

```
-----
```

```
-----
```

```
SQL> EXECUTE DBMS_RESULT_CACHE.MEMORY_REPORT;
```

```
Result Cache Memory Report
```

```
[Parameters]
```

```
Block Size = 1K bytes
```

```
Maximum Cache Size = 1376K bytes (1376 blocks)
```

```
Maximum Result Size = 68K bytes (68 blocks)
```

```
[Memory]
```

```
Total Memory = 5132 bytes [0.004% of the Shared Pool]
```

```
... Fixed Memory = 5132 bytes [0.004% of the Shared Pool]
```

```
... Dynamic Memory = 0 bytes [0.000% of the Shared Pool]
```

Why is the result cache not used? (Choose all that apply.)

- A. because the query uses SYSDATE
- B. because the query uses an alias for a column
- C. because the query uses the SLNOQ.CURRVAL sequence
- D. because the table might have an index on the SL_DATE column

Answer: A,C

Question No : 16

Identify the reason why chained or migrated rows in a table cause performance

degradation when you access them.

- A. because the ROWID changes for these rows
- B. because these rows make existing indexes invalid
- C. because it is not possible to use indexes to retrieve these rows
- D. because retrieving these rows requires more than one data block to be accessed
- E. because the rows force queries to use temporary segments while retrieving them

Answer: D

Question No : 17

Which three statements are true about performance analysis by SQL Performance Analyzer? (Choose three.)

- A. It detects changes in SQL execution plans.
- B. It shows only the overall impact on workload.
- C. It produces results that can be used to create the SQL plan baseline.
- D. It generates recommendations to run SQL Tuning Advisor to tune regressed SQL statements.
- E. The importance of SQL statements to be analyzed is based on the size of the objects accessed.

Answer: A,C,D

Question No : 18

During a proactive database performance monitoring routine, on examining the AWR report you find that log file sync appears among the top 5 wait events. What does this event indicate?

- A. Frequent logfile switches are occurring.
- B. Redo is generated faster than LGWR can write it out.
- C. Frequent commits or rollbacks are taking place in the application.
- D. Frequent incremental checkpoints are taking place in the database.

Answer: C

Question No : 19

You are asked to implement Automatic Memory Management (AMM) and you added two initialization parameters to do this. The following are the some of the parameter settings:

MEMORY_MAX_TARGET=600M

SGA_MAX_SIZE=500M

MEMORY_TARGET=600M

OPEN_CURSORS=300

SGA_TARGET=300M

PROCESSES=150

STATISTICS_LEVEL=BASIC

PGA_AGGREGATE_TARGET=0

When you tried to start the database instance with these parameter settings, you received the following error:

```
SQL> startup
```

```
ORA-00824: cannot set SGA_TARGET or MEMORY_TARGET due to existing internal settings, see alert log for more information
```

You modified the parameter settings as shown below and the database instance started successfully.

MEMORY_MAX_TARGET=600M

SGA_MAX_SIZE=400M

MEMORY_TARGET=550M

OPEN_CURSORS=400

SGA_TARGET=300M

PROCESSES=200

STATISTICS_LEVEL=TYPICAL

PGA_AGGREGATE_TARGET=100M

Analyze the parameter settings and identify why the instance failed to start the first time?

- A. because the STATISTICS_LEVEL parameter was set to BASIC
- B. because the PGA_AGGREGATE_TARGET parameter was set to zero
- C. because both the SGA_TARGET and MEMORY_TARGET parameters were set
- D. because MEMORY_TARGET and MEMORY_MAX_TARGET parameters were equal
- E. because the PROCESSES parameter was set to a value less than the minimum required

Answer: A

Question No : 20

The database application developers are planning to make some major schema changes such as creating new indexes and materialized views. They want to check the net impact of these changes on the workload performance. This activity has to be performed in the production database, so they want only the query part of the data manipulation language (DML) statements to be executed so that the side effects to the database or user data can be prevented. What should they use to achieve this?

- A. Database Replay
- B. SQL Tuning Advisor
- C. SQL Access Advisor
- D. SQL Performance Analyzer

Answer: D

Question No : 21

View the Exhibit and analyze a portion of the Active Session History (ASH) report marked with a red rectangle. Which two can be inferred from the output? (Choose two.)

Activity Over Time

- Analysis period is divided into smaller time slots
- Top 3 events are reported in each of those slots
- 'Slot Count' shows the number of ASH samples in that slot
- 'Event Count' shows the number of ASH samples waiting for that event in that slot
- '% Event' is 'Event Count' over all ASH samples in the analysis period

Slot Time (Duration)	Slot Count	Event	Event Count	% Event
17:13:00 (1.0 min)	6	CPU + Wait for CPU	6	0.25
17:14:00 (1.0 min)	634	buffer busy waits	590	24.67
		CPU + Wait for CPU	32	1.34
		log file sync	7	0.29
17:15:00 (1.0 min)	1,752	buffer busy waits	1,623	67.85
		CPU + Wait for CPU	61	2.55
		cursor: pin S	46	1.92

- A. increase in the number of active sessions
 B. unable to grow the database buffer cache as per demand
 C. increase in number of COMMITs and ROLLBACKS per second
 D. increase in the number of processes waiting for the buffer busy wait event

Answer: A,D

Question No : 22

Which two symptoms related to the database buffer cache together indicate that there are many full table scans happening? (Choose two.)

- A. too many buffer busy waits
 B. very low buffer cache-hit ratio
 C. very high buffer cache-hit ratio
 D. many waits on the db file scattered read event

Answer: B,D

Question No : 23

Your database supports many applications running on the middle tier. Many applications users create jobs for which you want to statistically measure workload as a part of performance management. What would you do to accomplish the task?

- A. Assign resource consumer group to jobs while creating the jobs.
- B. Create services for the applications and create jobs by using the DBMS_JOBS PL/SQL package.
- C. Query v\$SESSION to gather statistics of the individual sessions for the workload created by the jobs.
- D. Create services for the applications, create job class associated with the service, and then create jobs by using the job classes.

Answer: D

Question No : 24

View the Exhibit that shows your investigation on degraded performance. What would you do? (Choose all that apply.)

```
SQL> SELECT event, total_waits, total_timeouts, time_waited, average_wait
  2 FROM v$system_event
  3 WHERE event='buffer busy waits';
```

EVENT	TOTAL_WAITS	TOTAL_TIMEOUTS	TIME_WAITED	AVERAGE_WAIT
buffer busy waits	636528	1557	549700	.863591232

```
SQL> SELECT * FROM v$waitstat WHERE class='segment header';
```

CLASS	COUNT	TIME
segment header	61113	70278

- A. Shrink the temporary segment.
- B. Keep the large tables in the KEEP pool.
- C. Use free list groups wherever applicable.
- D. Increase the value set to the PGA_AGGREGATE_TARGET parameter.
- E. Switch the segments to Automatic Segment Space Management (ASSM).

Answer: C,E

Question No : 25

You are working as a DBA for a DSS system, on which applications primarily perform large sequential I/Os. The company wants to upgrade the storage used by the database to faster hard disks. What other factors would you consider before giving the go-ahead for the new hard disks? (Choose all that apply.)

- A. the sizing of PGA to meet the faster data access
- B. the sizing of buffer cache to meet the faster data access
- C. the time it takes for a submitted I/O request to be serviced by the storage
- D. the capacity of the I/O channel that connects the server to the storage array

Answer: C,D

Question No : 26

The OPTIMIZER_CAPTURE_SQL_PLAN_BASELINES parameter is set to TRUE for a 24x7 database instance. The space consumption has significantly increased by the automatically captured plan baselines. You are suggested to run the following block:

```
BEGIN
```

```
dbms_spm.configure('SPACE_BUDGET_PERCENT', 30);
```

```
END;
```

After a few weeks of work, you observe warnings on exceeding space usage by SQL Management Base (SMB) in alert log file. Which two actions can be solutions for this? (Choose two.)

- A. increasing the AWR retention period
- B. purging the older SQL management objects
- C. increasing the size of the SYSTEM tablespace
- D. increasing the size of the SYSAUX tablespace

Answer: B,D

Question No : 27

Identify the type of recommendation that can automatically be implemented by the Automatic Tuning Advisor as part of automatic SQL tuning task?

- A. recommendation for SQL profiles
- B. recommendation about effective indexing
- C. recommendation about stale or no statistics
- D. recommendation about the structure of SQL statements

Answer: A

Question No : 28

A user in a session executed the following SQL statement to set the optimizer mode:

```
ALTER SESSION SET OPTIMIZER_MODE = ALL_ROWS
```

What impact would it have on the goal of the optimizer for that session? (Choose all that apply.)

- A. Statement level OPTIMIZER_MODE hints take precedence over the session-level setting.
- B. The OPTIMIZER_MODE parameter set at instance level takes precedence over the session-level value.
- C. The optimizer uses a cost-based approach, regardless of the presence of statistics; it optimizes with a goal of best response time.
- D. The optimizer uses a cost-based approach for all SQL statements in the session, regardless of the presence of statistics; it optimizes with a goal of best throughput.

Answer: A,D

Question No : 29

View the Exhibit and examine the content of the V\$SGA_RESIZE_OPS view.

Dumps with VCE and PDF (+Free VCE Software)

COMPONENT	OPER_TYPE	OPER_MODE	INITIAL_SIZE	TARGET_SIZE	FINAL_SIZE	STATUS
ASM Buffer Cache	STATIC		0	0	0	COMPLETE
RECYCLE buffer cache	STATIC		0	0	0	COMPLETE
KEEP buffer cache	STATIC		0	8388608	8388608	COMPLETE
shared pool	STATIC		0	146800640	146800640	COMPLETE
large pool	STATIC		0	4194304	4194304	COMPLETE
java pool	STATIC		0	12582912	12582912	COMPLETE
streams pool	STATIC		0	0	0	COMPLETE
Shared IO Pool	STATIC		0	0	0	COMPLETE
DEFAULT buffer cache	STATIC		0	8388608	8388608	COMPLETE
DEFAULT 2K buffer cache	STATIC		0	0	0	COMPLETE
DEFAULT 4K buffer cache	STATIC		0	0	0	COMPLETE
DEFAULT 8K buffer cache	STATIC		0	0	0	COMPLETE
DEFAULT 16K buffer cache	STATIC		0	0	0	COMPLETE
DEFAULT 32K buffer cache	STATIC		0	0	0	COMPLETE
shared pool	GROW	IMMEDIATE	146800640	150994944	150994944	PENDING
DEFAULT buffer cache	GROW	DEFERRED	8388608	12582912	12582912	COMPLETE
shared pool	SHRINK	DEFERRED	150994944	146800640	146800640	COMPLETE

The Shared Pool was shrunk when the Default Buffer Cache required more memory. Why did the Java Pool not shrink when the Default Buffer Cache grew?

- A. because it is not an auto-tuned component
- B. because it uses memory available outside the total SGA size
- C. because it always requires three granules and cannot be reduced
- D. because it has less free memory than the Database Buffer Cache
- E. because it cannot be shrunk below the value set to the JAVA_POOL_SIZE parameter

Answer: E

Question No : 30

You work as a consultant DBA for various clients. A performance issue in one of the online transaction processing (OLTP) systems is reported to you and you received the Automatic Workload Repository (AWR) report generated in the database. The main sections of the AWR report are shown in the Exhibits.

View the Exhibit named DBTIME and note % DB Time.

Top 5 Timed Foreground Events

Event	Waits	Time(s)	Avg wait (ms)	% DB time	Wait Class
DB CPU		108		94.00	
db file sequential read	315,309	5	0	4.62	User I/O
latch: shared pool	1,227	0	0	0.41	Concurrency
db file parallel read	186	0	1	0.16	User I/O
control file sequential read	3,935	0	0	0.05	System I/O

View the Exhibit named TIMEMODEL and note what has contributed to % DB Time.

Time Model Statistics

- Total time in database user-calls (DB Time): 114.8s
- Statistics including the word "background" measure background process time, and so do not contribute to the DB time statistic
- Ordered by % of DB time desc, Statistic name

Statistic Name	Time (s)	% of DB Time
sql execute elapsed time	111.41	97.02
DB CPU	107.95	94.00
parse time elapsed	85.03	74.05
hard parse elapsed time	73.55	64.05
PL/SQL execution elapsed time	5.83	5.08
PL/SQL compilation elapsed time	2.20	1.92
hard parse (sharing criteria) elapsed time	0.46	0.40
connection management call elapsed time	0.44	0.38
hard parse (bind mismatch) elapsed time	0.27	0.24
sequence load elapsed time	0.06	0.05
repeated bind elapsed time	0.02	0.01
DB time	114.83	
background elapsed time	2.90	
background cpu time	0.95	

View the Exhibit named EFFICIENCY and examine the various percentages shown.

Instance Efficiency Percentages (Target 100%)

Buffer Nowait %:	100.00	Redo NoWait %:	100.00
Buffer Hit %:	74.48	In-memory Sort %:	100.00
Library Hit %:	68.42	Soft Parse %:	34.91
Execute to Parse %:	3.01	Latch Hit %:	99.95
Parse CPU to Parse Elapsed %:	0.01	% Non-Parse CPU:	34.94

Identify the option that has the correct answers for the questions given below:

- 1) Which factor indicates the main problem?
- 2) What is the main problem?
- 3) What solution would you recommend?

A. 1) The Exhibit DBTIME shows that the DB CPU event consumed very high % DB Time.
 2) The CPU is very slow.
 3) Increase the number of processors.

B. 1) hard parse elapsed time has the majority time in parse time elapsed, which is shown in the TIMEMODEL Exhibit.
 2) This is due to inadequate Database Buffer Cache.

3) Increase the size of database buffer cache.

C. 1) sql execute elapsed time and parse time elapsed are consuming very high % DB Time in the TIMEMODEL Exhibit.

2) There are too many soft parses.

3) Investigate the HOLD_CURSOR parameter setting in the application and set it appropriately.

D. 1) % Non-Parse CPU is low and Soft Parse % is also low in the EFFICIENCY Exhibit. This shows that very few statements are found in the cache.

2) There are too many hard parses.

3) Investigate the CURSOR_SHARING parameter setting and set it appropriately.

Answer: D

Question No : 31

You identified poorly performing SQL by analyzing the V\$SQL and V\$SQLSTATS views. You want to investigate the plans for these SQL statements. Which two methods can you use to pull out the execution plan from the library cache for the already executed SQL? (Choose two.)

A. Query V\$SQL_PLAN to view the execution plan.

B. Query DBA_HIST_SQL_PLAN to view the execution plan.

C. Copy and paste the SQL text from the V\$SQL view and use EXPLAIN PLAN to generate the execution plan.

D. Use the dbms_xplan.display_cursor function with the SQL ID and child number to generate the execution plan.

Answer: A,D

Question No : 32

You observed very low cache-hit ratio in your database as shown below:

```
SQL> SELECT (1-((phy.value-phyd.value) / (cur.value + con.value))) * 100
```

```
2 "Cache Hit ratio"
```

```
3 FROM v$sysstat cur, v$sysstat con, v$sysstat phy, v$sysstat phyd
```

```
4 WHERE cur.name = 'db block gets'
```

5 AND con.name = 'consistent gets'

6 AND phy.name = 'physical reads'

7 AND phyd.name = 'physical reads direct';

Cache Hit Ratio

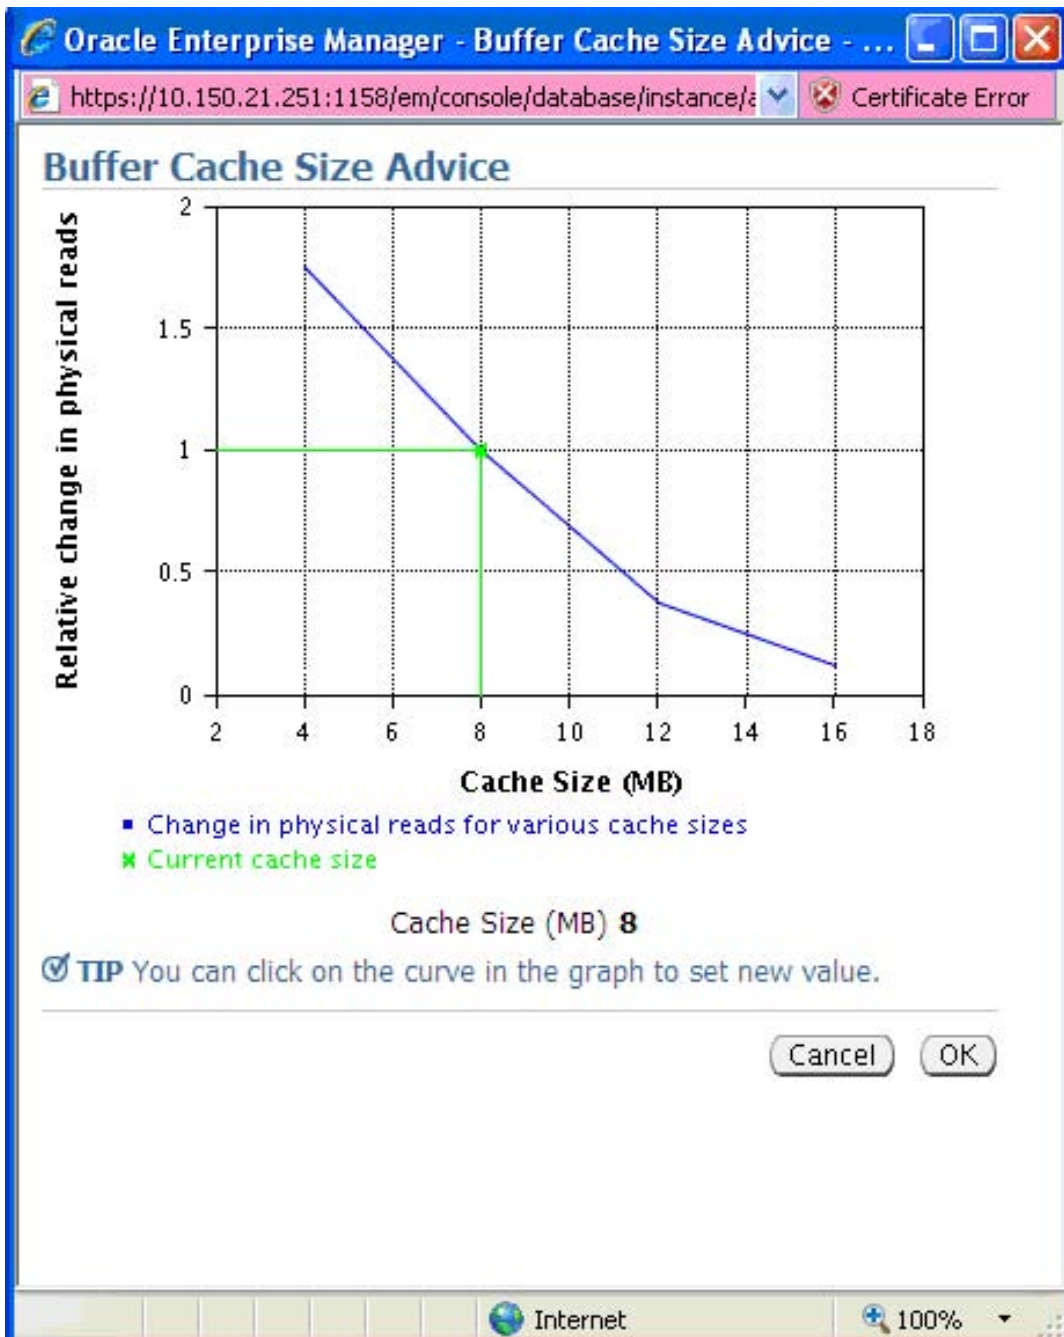
68.43

After further investigation, you decided to increase the database buffer cache size. You are using the

Buffer Cache Advisor to check the appropriate size for the buffer cache.

View the Exhibit and examine the graph shown by the advisory.

What can you infer from the graph?



- A. The buffer cache size can be set to 12 MB for optimal performance.
- B. The buffer cache size can be set to more than 16 MB to get more benefit.
- C. The buffer cache size can be set to more than 16 MB but there would be no benefit from it.
- D. The buffer cache size cannot be set to more than 16 MB because the SGA_MAX_SIZE value does not allow that.

Answer: B

Question No : 33

Examine the output of the following query:

```
SQL> SELECT c.name,a.addr,a.gets,a.misses,a.sleeps,  
2 a.immediate_gets,a.immediate_misses,b.pid  
3 FROM v$latch a, v$latchholder b, v$latchname c  
4 WHERE a.addr = b.laddr(+) and a.latch# = c.latch#  
5 AND c.name LIKE '&latch_name%' ORDER BY a.latch#;  
  
LATCH NAME ADDR GETS MISSES SLEEPS IMMEDIATE_GETS IMMEDIATE_MISSES  
-----  
  
shared pool 20016544 8520540 14112 3137 0 0
```

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

- A. The latch was requested in no wait mode.
- B. The latch was requested in willing-to-wait mode.
- C. The shared pool operations to allocate and free memory in it caused the misses.
- D. Automatic Memory Management is not enabled because the IMMEDIATE_GETS and IMMEDIATE_MISSES columns have zero in them.

Answer: B,C

Question No : 34

You identified poorly performing SQL statements that use bind variables. You want to examine the execution plans for the most recently executed statements for further diagnosis. Which method would you recommend to view the execution plans?

- A. SQL*Plus Autotrace
- B. the dbms.xplan package
- C. EXPLAIN PLAN for SQL statements
- D. the SQL trace utility to generate trace files and read the formatted output using the tkprof utility

Answer: B

Question No : 35

You queried V\$SYSTEM_EVENT in your online transaction processing (OLTP) environment to check the following wait events:

Event Total Seconds Total Average Wait

Name Waits Waiting Timeouts (in secs)

direct path read temp 18,274 11,312 16,282 .619

direct path write temp 3,251 416 25 .127

What do these wait events indicate?

- A. The DBWn processes are not creating free buffers fast enough.
- B. The buffer cache is too small and block replacement is excessive.
- C. The aggressive checkpointing policy is causing too many writes to disk.
- D. The work areas are too large to fit in memory and are, therefore, written to disk.

Answer: D

Question No : 36

You observed in your database that the queries formerly executed in 5 seconds now take longer time. View the Exhibit and analyze the output of the query that you executed to find the problem. Which is the primary concern?

Dumps with VCE and PDF (+Free VCE Software)

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

NAME	VALUE
physical reads direct temporary tablespace	502
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	227752
table fetch by rowid	220813
table fetch continued row	132046
table lookup prefetch client count	0
LOB table id lookup cache misses	0

18 rows selected.

- A. the number of disk sorts
- B. the number of full table scans
- C. the number of chained and migrated rows
- D. the number of rows fetched based on rowid

Answer: C

Question No : 37

View the Exhibit and examine the query based on the DBA_THRESHOLDS view and its partial output :

Which observation about the output is correct?

```
SQL>SELECT metrics_name,warning_operator OPR,
       warning_value WARN_VAL,observation_period OBS_PER,
       consecutive_occurrences CONS_OCC, object_type OBJ_TYPE,
       object_name OBJ_NAME
FROM dba_thresholds;
```

METRICS_NAME	OPR	WARN_VAL	OBS_PER	CONS_OCC	OBJ_TYPE	OBJ_NAME
Average Users Waiting Counts	GT	30	1	3	EVENT_CLASS	Commit
Blocked User Session Count	GT	0	1	15	SESSION	
Database Time Spent Waiting (%)	GT	30	1	3	EVENT_CLASS	Concurrency
Logons Per Sec	GE	100	1	2	SYSTEM	

- A. Logons Per Sec triggers an alert when the number of users logging on in an interval of one second is equal to or greater than 100.
- B. Blocked User Session Count triggers an alert based on the

ELAPSED_TIME_PER_CALL and DB_TIME_WAITING metrics collection per minute.

C. Database Time Spent Waiting(%) triggers an alert when more than 30% of the DB time per second for three consecutive seconds is spent on lock waits.

D. Average Users Waiting Counts triggers an alert when the count is more than 30 for 3 consecutive 1 minute intervals and is related to the log file sync wait event.

Answer: D

Question No : 38

A batch workload that historically completed in the maintenance window between 10:00 PM and midnight is currently showing poor performance and completing at 2 AM. To help in the diagnosis of the performance degradation, the senior DBA in your organization asks you to execute the awrddrpt.sql script to produce a Compare Periods report. Which two statements are true about the report produced by this script? (Choose two.)

A. It is refreshed automatically based on the moving window baseline.

B. It compares details between any two selected time periods of the same duration.

C. It normalizes the statistics by the amount of time spent on the database for each time period.

D. It compares details between two consecutive time periods of the same or different durations and is refreshed every 60 minutes.

Answer: B,C

Question No : 39

View Exhibit1 and examine a portion of Active Session History(ASH) report.

Top DB Objects

- With respect to Application, Cluster, User I/O and buffer busy waits only.

Object ID	% Activity	Event	% Event	Object Name (Type)	Tablespace
77025	92.52	buffer busy waits	92.52	SPC.SPCT (TABLE)	TBSSPC

[Back to Top Objects/Files/Latches](#)

[Back to Top](#)

Top DB Files

No data exists for this section of the report.

[Back to Top Objects/Files/Latches](#)

[Back to Top](#)

Top Latches

No data exists for this section of the report.

[Back to Top Objects/Files/Latches](#)

[Back to Top](#)

Activity Over Time

- Analysis period is divided into smaller time slots
- Top 3 events are reported in each of those slots
- 'Slot Count' shows the number of ASH samples in that slot
- 'Event Count' shows the number of ASH samples waiting for that event in that slot
- '% Event' is 'Event Count' over all ASH samples in the analysis period

Slot Time (Duration)	Slot Count	Event	Event Count	% Event
17:13 00 (1.0 min)	6	CPU + Wait for CPU	6	0.25
17:14 00 (1.0 min)	634	buffer busy waits	590	24.67
		CPU + Wait for CPU	32	1.34
		log file sync	7	0.29
17:15 00 (1.0 min)	1,752	buffer busy waits	1,623	67.85
		CPU + Wait for CPU	61	2.55
		cursor: pin S	46	1.92

View Exhibit2 to examine the tablespaces.

TABLESPACE_NAME	STATUS	CONTENTS	EXTENT_MAN	ALLOCATION	SEGMENT
SYSTEM	ONLINE	PERMANENT	LOCAL	SYSTEM	MANUAL
SYSAUX	ONLINE	PERMANENT	LOCAL	SYSTEM	AUTO
UNDOTBS1	ONLINE	UNDO	LOCAL	SYSTEM	MANUAL
TEMP	ONLINE	TEMPORARY	LOCAL	UNIFORM	MANUAL
USERS	ONLINE	PERMANENT	LOCAL	SYSTEM	AUTO
EXAMPLE	ONLINE	PERMANENT	LOCAL	SYSTEM	AUTO
SOEINDEX	ONLINE	PERMANENT	LOCAL	UNIFORM	AUTO
SOE	ONLINE	PERMANENT	LOCAL	UNIFORM	AUTO
TBSSPC	ONLINE	PERMANENT	LOCAL	SYSTEM	MANUAL

Which action would you suggest after analyzing the output?

- A. Increase the log buffer size.
- B. Rebuild indexes on SPCT table.
- C. Decreasing the number of free lists in segment SPCT.
- D. Move SPCT table to a different tablespace that is locally managed with automatic segment space management.

Answer: D

Question No : 40

View the Exhibit and examine a portion of the output obtained from the following query:

STAT_ID	STAT_NAME	VALUE
3649082374	DB time	61021783
2748282437	DB CPU	3890625
4157170894	background elapsed time	42472524
2451517896	background cpu time	2796875
4127043053	sequence load elapsed time	0
1431595225	parse time elapsed	10983653
372226525	hard parse elapsed time	10480831
2821698184	sql execute elapsed time	50353110
1990024365	connection management call elapsed time	855906

SQL> SELECT * FROM v\$sys_time_model;

Select three correct interpretations of the time model statistics. (Choose three.)

- A. DB time includes the wait time of all the nonidle and idle user sessions.
- B. SQL execute elapsed time includes the time spent in performing fetches of query results.
- C. DB CPU includes the CPU time spent on database user-level calls and background CPU time.
- D. SQL execute elapsed time includes components of the hard parse elapsed time like bind elapsed time.
- E. DB time includes the connection management call elapsed time excluding the background process time.

Answer: B,D,E

Question No : 41

You are working on an online transaction processing (OLTP) system. The CUST table in the SH schema was populated by direct load and after that it has gone through a lot of updates and deletions. The statistics for the CUST and SALES tables were updated recently. View the Exhibit and examine the query plan. The query is performing a lot of I/O for a query that fetches only 168 rows. To investigate further, you queried the ALL_TABLES view to find out PCTUSED, PCTFREE, and the number of rows in the CUST table, as given below:

```
SQL> SELECT table_name ,blocks, pct_used, pct_free, num_rows
```

```
2 FROM all_tables
```

```
3* WHERE table_name = 'CUST';
```

```
TABLE_NAME BLOCKS PCT_USED PCT_FREE NUM_ROWS
```

```
-----
CUST 13768 80 111060
```

What would you recommend to reduce the I/O?

```
SQL> SELECT time_id, quantity_sold, amount_sold
FROM sales s, cust c
WHERE c.cust_id = s.cust_id AND cust_first_name = 'Jimmy'
ORDER BY time_id
```

```
168 rows selected.
```

```
Execution Plan
```

```
-----
Plan hash value: 980033780
```

Id	Operation	Name	Rows	Bytes	TempSpc	Cost (%CPU)	Time	Pstart	Pstop
0	SELECT STATEMENT		11120	358K		4344 (1)	00:00:53		
1	SORT ORDER BY		11120	358K	1064K	4344 (1)	00:00:53		
2	HASH JOIN		11120	358K		4243 (1)	00:00:51		
3	TABLE ACCESS FULL	CUST	85	1020		3738 (1)	00:00:45		
4	PARTITION RANGE ALL		918K	18M		498 (4)	00:00:06	1	28
5	TABLE ACCESS FULL	SALES	918K	18M		498 (4)	00:00:06	1	28

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

```
-----
2 - access("C"."CUST_ID"="S"."CUST_ID")
3 - filter("CUST_FIRST_NAME"='Jimmy')
```

```
Statistics
```

```
-----
14863 recursive calls
0 db block gets
18448 consistent gets
13688 physical reads
0 redo size
4716 bytes sent via SQL*Net to client
541 bytes received via SQL*Net from client
13 SQL*Net roundtrips to/from client
137 sorts (memory)
0 sorts (disk)
168 rows processed
```


- A. reorganizing the table to use fewer blocks
- B. rebuilding the index on the CUST_FIRST_NAME column
- C. increasing the value for the PCTFREE attribute for the CUST table
- D. increasing the value for the PGA_AGGREGATE_TARGET initialization parameter

Answer: A

Question No : 42

You have the automatic optimizer statistics gathering enabled for a 24x7 database as part of the automatic maintenance task. Because the automatic optimizer statistics collection runs during maintenance window(in the night), the statistics on tables that are significantly modified throughout the day becomes stale. This leads to poor query performance on these tables and you want to minimize the

overhead in statistics collection. Identify a solution to this problem.

- A. Change the STALE_PERCENT preference to zero for the tables.
- B. Set the OPTIMIZER_USE_PENDING parameter statistics to TRUE.
- C. Set the OPTIMIZER_DYNAMIC_SAMPLING parameter to zero during day time.
- D. Gather statistics on the table when the tables have a representative number of rows, and then lock the statistics.

Answer: D

Question No : 43

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

Baseline Metric Thresholds

Threshold Configuration		Metric Analysis			
Quick Configuration					
View	Metrics with Thresholds				
Category/Name	Alerts (Last 7 Days)	AWR Baseline	Threshold Type	Adaptive	Edit
▼ Workload Volume Metrics	0 0				
Cumulative Logons (per second)	0 0		Fixed Values		
Current Open Cursors Count	0 0		Fixed Values		
I/O Requests (per second)	0 0	SYSTEM_MOVING_WINDOW	Significance Level	✓	
Number of Transactions (per second)	0 0		Fixed Values		
User Calls (per second)	0 0	SYSTEM_MOVING_WINDOW	Significance Level	✓	

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

Answer: A

Question No : 44

You work as a DBA for a company and you have the responsibility of managing one of its online transaction processing (OLTP) systems. The database encountered performance-related problems and you generated an Automatic Workload Repository (AWR) report to investigate it further. View the Exhibits and examine the AWR report.

Top 5 Timed Foreground Events

Event	Waits	Time(s)	Avg wait (ms)	% DB time	Wait Class
DB CPU		584		29.08	
library cache: mutex X	14,721	71	5	3.53	Concurrency
latch: shared pool	1,158	55	48	2.76	Concurrency
cursor: pin S wait on X	3,777	50	13	2.50	Concurrency
log file sync	672	17	25	0.83	Commit

Time Model Statistics

- Total time in database user-calls (DB Time): 2008.5s
- Statistics including the word "background" measure background process time, and so do not contribute to the DB time statistic
- Ordered by % of DB time desc, Statistic name

Statistic Name	Time (s)	% of DB Time
sql execute elapsed time	1,731.94	86.23
DB CPU	584.11	29.08
parse time elapsed	533.72	26.57
hard parse elapsed time	416.43	20.73
connection management call elapsed time	33.26	1.66
PL/SQL compilation elapsed time	10.58	0.53
Java execution elapsed time	8.01	0.40
failed parse elapsed time	5.20	0.26
PL/SQL execution elapsed time	3.66	0.18
hard parse (sharing criteria) elapsed time	1.94	0.10
hard parse (bind mismatch) elapsed time	1.33	0.07
sequence load elapsed time	0.41	0.02
repeated bind elapsed time	0.05	0.00
DB time	2,008.48	
background elapsed time	32.06	
background cpu time	4.79	

Load Profile

	Per Second	Per Transaction	Per Exec	Per Call
DB Time(s):	3.8	12.6	0.01	0.00
DB CPU(s):	1.1	3.7	0.00	0.00
Redo size:	6,062.3	20,190.1		
Logical reads:	5,982.5	19,924.3		
Block changes:	25.5	84.9		
Physical reads:	2,778.2	9,252.7		
Physical writes:	2.9	9.7		
User calls:	1,263.4	4,207.7		
Parses:	506.6	1,687.3		
Hard parses:	53.3	177.5		
W/A MB processed:	726,646.9	2,420,040.5		
Logons:	1.1	3.5		
Executes:	513.1	1,708.9		
Rollbacks:	0.1	0.3		
Transactions:	0.3			

Dictionary Cache Stats

- "Pct Misses" should be very low (< 2% in most cases)
- "Final Usage" is the number of cache entries being used

Cache	Get Requests	Pct Miss	Scan Reqs	Pct Miss	Mod Reqs	Final Usage
do_awr_control	13	69.23	0		2	1
do_database_links	1,074	0.56	0		0	0
do_global_oids	15,419	2.67	0		0	13
do_histogram_data	77,565	21.21	0		0	571
do_histogram_defs	168,045	23.16	0		0	1,014
do_object_grants	44,042	4.17	0		0	59
do_objects	358,789	3.30	0		0	396
do_profiles	548	2.19	0		0	1
do_rollback_segments	230	0.00	0		0	38
do_segments	99,805	15.72	0		5	279
do_sequences	25	100.00	0		25	0
do_tablespaces	85,868	0.04	0		0	5
do_users	179,387	0.35	0		0	20
global database name	927	0.11	0		0	1
kqsubheap_object	197	30.46	0		0	0
outstanding_alerts	15	84.74	0		0	1

[Back to Top](#)

Library Cache Activity

- "Pct Misses" should be very low

Namespace	Get Requests	Pct Miss	Pin Requests	Pct Miss	Reloads	Invali- dations
BODY	1,832	1.36	3,673	1.55	23	0
CLUSTER	2,761	1.81	1,590	3.14	0	0
INDEX	947	35.59	947	35.80	1	0
JAVA DATA	4	75.00	873	0.69	0	0
SQL AREA	340,330	23.79	602,683	12.78	22,142	5,231
TABLE/PROCEDURE	145,469	2.49	191,059	8.55	5,812	0
TRIGGER	5,539	0.23	5,539	0.29	0	0

What could be the problem in this database?

- Java pool is not configured.
- The CPU in the system is slow.
- The shared pool size is inadequate.
- The database buffer cache is inadequate.
- The OPEN_CURSORS parameter is set to a small value.

Answer: C

Question No : 45

After running SQL Performance Analyzer (SPA), you observe a few regressed SQL statements in the SPA output. Identify the two actions that you would suggest for these regressed SQL statements. (Choose two.)

- A. Running SQL Access Advisor
- B. Adding them to SQL Plan Baseline
- C. Submitting them to SQL Tuning Advisor
- D. Running Automatic Database Diagnostic Monitor (ADDM)

Answer: B,C

Question No : 46

The Automatic Workload Repository (AWR) snapshot interval is set to 1 hour. A user of your database notifies you that the system was slow between 10:00 PM and 10:10 PM, but the Automatic Database Diagnostic Management (ADDM) analysis for the time period between 10:00 PM and 11:00 PM does not show this performance problem. What would you suggest to analyze this performance problem for the user?

- A. SQL Tuning Advisor
- B. SQL Performance Analyzer
- C. Active Session History report
- D. AWR Compare Periods report

Answer: C

Question No : 47

On examining the Automatic Workload Repository (AWR) report, you find that log file sync appears among the top 5 wait events. On further investigation, you notice that the average time waited for this event is low, but the number of waits are high. What does this imply?

- A. DBWR is slow.
- B. LGWR is slow due to slow I/O.
- C. Redo log buffer size is insufficient.
- D. The application might be committing frequently.

Answer: D

Question No : 48

View the Exhibit and examine the plans in the SQL plan baseline. You observe that the last two SELECT statements (outlined in red and blue) are identical, but have two different plans. Also, the second plan (in blue) is not being used by the optimizer. Which two steps do you perform to ensure that the optimizer uses the second plan (in blue) if the plan is more efficient? (Choose two.)

Select Name	SQL Text	Enabled	Accepted	Fixed	Auto Purge	Created	Last Modified
<input type="checkbox"/> SYS_SQL_PLAN_ba5e158c52d2775d	DELETE FROM PLAN_TABLE WHERE STATEMENT_ID=:1	YES	YES	NO	YES	Jul 20, 2008 5:12:07 PM	Jul 20, 2008 5:12:07 PM
<input type="checkbox"/> SYS_SQL_PLAN_331fa51fdf463620	SELECT PLAN_TABLE_OUTPUT FROM TABLE(DBMS_XPLAN.DIS...	YES	YES	NO	YES	Jul 20, 2008 5:12:06 PM	Jul 20, 2008 5:12:06 PM
<input type="checkbox"/> SYS_SQL_PLAN_184921b7cf314e9e	select * from employees where job_id='CLERK'	YES	YES	NO	YES	Jul 20, 2008 5:15:31 PM	Jul 20, 2008 5:15:31 PM
<input type="checkbox"/> SYS_SQL_PLAN_184921b70572d2e8	select * from employees where job_id='CLERK'	YES	NO	NO	YES	Jul 20, 2008 5:17:33 PM	Jul 20, 2008 5:17:33 PM

✓ TIP The table will display maximum of 2000 rows. Use search criteria to get the desired results.

- A. Mark the second plan (in blue) as fixed.
- B. Use SQL Tuning Advisor to tune the plan.
- C. Use SQL Access Advisor to tune the plan.
- D. Use the DBMS_SPM.EVOLVE_SQL_PLAN_BASELINE function to evolve the plan.

Answer: B,D

Question No : 49

The users in your online transaction processing (OLTP) environment complain that the query response time has increased considerably. You checked the Program Global Area (PGA)- related information in the AWR report to investigate it further.

View the Exhibit named PGAHIT and note the PGA cache-hit percentage.

PGA Aggr Summary

- PGA cache hit % - percentage of W/A (WorkArea) data processed only in-memory

PGA Cache Hit %	W/A MB Processed	Extra W/A MB Read/Written
25.99	15,966	45,458

View the Exhibit named PGAADV and analyze the content.

PGA Memory Advisory

- When using Auto Memory Mgmt, minimally choose a pga_aggregate_target value where Estd PGA Overalloc Count is 0

PGA Target Est (MB)	Size Factr	W/A MB Processed	Estd Extra W/A MB Read/ /written to Disk	Estd PGA Cache Hit %	Estd PGA Overalloc Count	Estd Time
10	0.50	15,972.74	32,222.39	33.00	854	504,221
15	0.75	15,972.74	32,222.39	33.00	854	504,221
20	1.00	15,972.74	22,728.29	41.00	854	404,893
24	1.20	15,972.74	22,728.29	41.00	854	404,893
28	1.40	15,972.74	22,728.29	41.00	854	404,893
32	1.60	15,972.74	22,728.29	41.00	854	404,893
36	1.80	15,972.74	22,728.29	41.00	854	404,893
40	2.00	15,972.74	22,728.29	41.00	854	404,893
60	3.00	15,972.74	22,728.29	41.00	854	404,893
80	4.00	15,972.74	22,728.29	41.00	854	404,893
120	6.00	15,972.74	22,687.95	41.00	346	404,471
160	8.00	15,972.74	7,324.74	69.00	0	243,740

If you want to increase the size of PGA, you have sufficient memory available in the system. What would be the correct decision?

- A. Double the size of PGA.
- B. Increase the size of PGA to 120 MB.
- C. Increase the size of PGA to 160 MB or more.
- D. Leave the PGA at its current size and investigate other possible causes.

Answer: C

Question No : 50

Examine the following information from the AWR report:

Top 5 Timed Events

~~~~~ % Total

Event Waits Time (s) Call Time

-----

CPU time 559 88.80

log file parallel write 2,181 28 4.42

SQL\*Net more data from client 516,611 27 4.24

db file parallel write 13,383 13 2.04

db file sequential read 563 2 .27

Which could cause waits on the db file sequential read event?

- A. too many full table scans
- B. too large a database buffer cache
- C. too many waits occurring while reading multiple blocks synchronously
- D. too many index range scans occurring which refers to blocks not in the database buffer cache

**Answer: D**

### **Question No : 51**

You work as a DBA in ABC Corp. You are working on an online transaction processing (OLTP) database. The database size has grown to 100 GB. The database will support new applications and more users will be added. Presently, the database uses normal file system storage. Before the new applications are added, you want to move the storage to RAID. You are in consultation with the OS Administrator for finalizing the RAID option for the storage.

You have the following requirements for the database:

- 1) Less mean time to recover (MTTR)
- 2) Striped for performance
- 3) Mirrored for safety

Which RAID option would you recommend?

- A. RAID level 0
- B. RAID level 1
- C. RAID level 0 over 1(1+0)
- D. RAID level 1 over 0(0+1)

**Answer: C**

### **Question No : 52**

You work as a DBA and have the responsibility of managing a large online transaction processing (OLTP) system. You used three queries to check the database performance as

shown in the Exhibit. View the Exhibit and analyze the output.

```
SQL> SELECT (1-((phy.value-phyd.value) / (cur.value + con.value))) * 100
 2  "Cache Hit ratio"
 3  FROM v$sysstat cur, v$sysstat con, v$sysstat phy, v$sysstat phyd
 4  WHERE cur.name = 'db block gets'
 5  AND con.name = 'consistent gets'
 6  AND phy.name = 'physical reads'
 7  AND phyd.name = 'physical reads direct';
```

```
Cache Hit Ratio
-----
          99.43
```

```
SQL> SELECT event, total_waits, total_timeouts, time_waited, average_wait
 2  FROM v$system_event
 3  WHERE event='buffer busy waits';
```

| EVENT             | TOTAL_WAITS | TOTAL_TIMEOUTS | TIME_WAITED | AVERAGE_WAIT |
|-------------------|-------------|----------------|-------------|--------------|
| buffer busy waits | 36528       | 1557           | 549700      | 15.04872974  |

```
SQL> SELECT * FROM v$waitstat WHERE class='data block';
```

| CLASS      | COUNT   | TIME    |
|------------|---------|---------|
| data block | 1961113 | 1870278 |

What conclusion can you draw from this?

- A. There are many physical I/Os happening.
- B. There are many full table scans happening.
- C. The data blocks are aging out of the buffer cache very fast.
- D. There are many cursors trying to access the same data blocks.
- E. The DBWn processes are not freeing sufficient buffers to meet the demand.

**Answer: D**

### Question No : 53

Your company wants to migrate a single instance Oracle 11g Release 1 database to a RAC environment. The database supports middle-tier applications using connection pooling for connecting to database.

Being a part of the performance management team, you want to test the performance of the SQL statements and the peak workload of the applications in the new environment before upgrading the production system.

Given below are some of the steps in random order to accomplish the task:

- 1) Capture the real workload on the existing system.
- 2) Create a SQL Tuning Set (STS) for the statements executed by the application.
- 3) Transport STS on the test environment and use SQL Tuning Advisor to compare the before and after performance.
- 4) Replay the workload on the test environment.
- 5) Analyze the divergence after replay.
- 6) Use Automatic Workload Repository (AWR) to capture the information for top SQL on the existing system.
- 7) Transport STS on the test environment and use SQL Performance Analyzer to compare the before and after performance.

Identify the correct steps required to perform the test.

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

**Answer: C**

**Question No : 54**

View the Exhibit and examine the partial output from the query on the v\$file\_histogram view: Which statement correctly interprets the output displayed in the Exhibit?

```
SQL> SELECT * FROM v$file_histogram;
```

| FILE# | SINGLEBLKRDTIM_MILLI | SINGLEBLKRDS |
|-------|----------------------|--------------|
| 1     | 1                    | 49           |
| 1     | 2                    | 58           |
| 1     | 4                    | 138          |
| 1     | 8                    | 498          |
| 1     | 16                   | 1353         |
| 1     | 32                   | 1499         |
| 1     | 64                   | 1734         |
| 1     | 128                  | 576          |
| 1     | 256                  | 56           |
| 1     | 512                  | 51           |

- A. The database waited 58 times in a time interval of two milliseconds for FILE# 1.
- B. The database experienced 58 waits of two milliseconds duration each for FILE# 1.
- C. The database performed 58 synchronous single block reads in a time interval of two milliseconds for FILE# 1.
- D. The database experienced 58 waits of duration less than two milliseconds each that are not included in any smaller bucket.

**Answer: D**

**Question No : 55**

View the Exhibit. You want to set the threshold for the number of transactions (per second) metric with the threshold type as significance level. Which baseline would you choose to make the threshold adaptive?

| Select                           | Name                      | Type                   | Valid | Statistics Computed | Last Time Computed       | Start Time               | End Time                |
|----------------------------------|---------------------------|------------------------|-------|---------------------|--------------------------|--------------------------|-------------------------|
| <input checked="" type="radio"/> | LAST_NIGHT                | STATIC                 | Yes   | Yes                 | Jun 19, 2008 12:26:53 PM | Jun 15, 2008 10:00:06 AM | Jun 19, 2008 8:00:50 AM |
| <input type="radio"/>            | MONDAY_1_2008-06-24.02:00 | GENERATED              | No    | No                  |                          | Jun 26, 2008 1:57:52 PM  | Jun 26, 2008 1:57:52 PM |
| <input type="radio"/>            | MONDAY_1_2008-07-01.02:00 | GENERATED              | Yes   | No                  |                          | Jun 30, 2008 10:00:06 PM | Jul 1, 2008 2:00:13 AM  |
| <input type="radio"/>            | MONDAY_1_2008-07-08.02:00 | GENERATED              | Yes   | No                  |                          | Jul 7, 2008 10:00:38 PM  | Jul 8, 2008 2:00:44 AM  |
| <input type="radio"/>            | MONDAY_1_2008-07-15.02:00 | GENERATED              | Yes   | No                  |                          | Jul 14, 2008 10:00:06 PM | Jul 15, 2008 2:00:13 AM |
| <input type="radio"/>            | SYSTEM_MOVING_WINDOW      | MOVING_WINDOW (8 Days) | Yes   | Yes                 | Jul 13, 2008 2:00:00 PM  | Jul 7, 2008 7:00:33 PM   | Jul 15, 2008 6:00:40 PM |
| <input type="radio"/>            | THU_SECTION               | STATIC                 | Yes   | Yes                 | Jul 15, 2008 4:12:14 PM  | Jul 9, 2008 4:00:45 PM   | Jul 10, 2008 2:00:20 PM |

- A. LAST\_NIGHT
- B. THU\_SECTION
- C. SYSTEM\_MOVING\_WINDOW
- D. MONDAY\_1\_2008-07-15.02:00

**Answer: C**

**Question No : 56**

You work as a DBA for a company and as a performance improvement measure, you implemented the result cache in your database.

View the Exhibit named SETTING and note the result cache settings.

```
SQL> SHOW PARAMETER RESULT
```

| NAME                           | TYPE        | VALUE |
|--------------------------------|-------------|-------|
| client_result_cache_lag        | big integer | 3000  |
| client_result_cache_size       | big integer | 0     |
| result_cache_max_result        | integer     | 5     |
| result_cache_max_size          | big integer | 1376K |
| result_cache_mode              | string      | FORCE |
| result_cache_remote_expiration | integer     | 0     |

Many users in the company state that performance has improved on the queries they use but some users complain that they have not got any performance benefit on the queries they use.

You checked all the queries they use and the following is one of them:

```
SQL> SELECT slnoq.currval as "SLNO", prod_id, pdname,  
2 cust_name FROM sales WHERE sl_date < sysdate;
```

View the Exhibit named TEST and examine the testing performed to check this.



```
SQL> EXECUTE DBMS_RESULT_CACHE.FLUSH;
```

PL/SQL procedure successfully completed.

```
SQL> EXECUTE DBMS_RESULT_CACHE.MEMORY_REPORT;
```

```
Result Cache Memory Report
```

```
[Parameters]
```

```
Block Size = 1K bytes
```

```
Maximum Cache Size = 1376K bytes (1376 blocks)
```

```
Maximum Result Size = 68K bytes (68 blocks)
```

```
[Memory]
```

```
Total Memory = 5132 bytes [0.004% of the Shared Pool]
```

```
... Fixed Memory = 5132 bytes [0.004% of the Shared Pool]
```

```
... Dynamic Memory = 0 bytes [0.000% of the Shared Pool]
```

PL/SQL procedure successfully completed.

```
SQL> SELECT slnoq.currval as "SLNO", prod_id, pname, cust_name
```

```
2 FROM sales
```

```
3 WHERE sl_date < sysdate;
```

```
-----
```

```
-----
```

```
SQL> EXECUTE DBMS_RESULT_CACHE.MEMORY_REPORT;
```

```
Result Cache Memory Report
```

```
[Parameters]
```

```
Block Size = 1K bytes
```

```
Maximum Cache Size = 1376K bytes (1376 blocks)
```

```
Maximum Result Size = 68K bytes (68 blocks)
```

```
[Memory]
```

```
Total Memory = 5132 bytes [0.004% of the Shared Pool]
```

```
... Fixed Memory = 5132 bytes [0.004% of the Shared Pool]
```

```
... Dynamic Memory = 0 bytes [0.000% of the Shared Pool]
```

1

Why is the result cache not used? (Choose all that apply.)

- A. because the query uses SYSDATE
- B. because the query uses an alias for a column
- C. because the query uses the SLNOQ.CURRVAL sequence
- D. because the CLIENT\_RESULT\_CACHE\_SIZE parameter is set to 0

**Answer: A,C**

### Question No : 57

You are working on a database that was upgraded to Oracle Database 11g from Oracle Database 9i. An ADDM finding in this database says that the shared pool is inadequately