Exam Name:	**Oracle 9i Database: Performance Tuning**		
Exam Type:	**Oracle**		
Exam Code:	**1Z0-033**	Total Questions:	**165**

Question: 1
Which statements are true regarding configuring multiple buffer pools in the Database Buffer Cache that is using only the standard block size? (Choose all that apply.)

A. Both the KEEP and RECYCLE pools must be configured.
B. The RECYCLE pool needs to be configured before dropping an object.
C. The DEFAULT pool needs to be removed if multiple buffer pools are being configured.
D. Either the KEEP pool or the RECYCLE pool can be configured along with the DEFAULT pool.

Answer: D

Question: 2
Smith is a DBA with ABC Corp. He is in the process of tuning the database. He has received various problem statements from the customers. Identify three appropriately defined problem statements that Smith can use to tune the database correctly. (Choose three.)

A. The system is too slow.
B. Online response is very slow.
C. Sometimes the batch process throws errors.
D. The weekly backup took 30 minutes, and earlier it used to take 5 minutes.
E. The batch process is taking 2 hours, whereas it should take only 30 minutes.
F. The system currently has 100 users. We wish to increase the number of users to 250. The users' response time should not be affected.

Answer: D, E, F

Question: 3
You execute the following query to check resource consumption in the current plan used by the database:

SQL> SELECT name,active_sessions,requests,queue_length,current_undo_consumption "UNDO
CONSUMPTION" FROM V$rsrc_consumer_group;
NAME ACTIVE_SESSIONS REQUESTS QUEUE_LENGTH UNDO CONSUMPTION
--------------- ------------------------ -------------- -------------------------------- -------------------
OTHER_GROUPS 1 1 0 0
GROUP1 2 4 2 0

Which statement is correct when a user belonging to GROUP1 tries to start a session?

A. The user can login but the session will be queued.
B. The user's attempt to start a session fails with an error.
C. The session starts successfully and allows all the operations.
D. The session starts successfully but cannot perform data definition language (DDL) operations.
E. The session opens successfully but cannot perform data manipulation language (DML) operations.

Answer: A

Question: 4
You require that a warning is written to a trace file without impacting the application when the session's Java state exceeds 2GB. Which parameter would you modify?

A. JAVA_MAX_SESSIONSPACE_SIZE
B. JAVA_SOFT_SESSIONSPACE_LIMIT

C. SESSIONS and JAVA_SOFT_SESSIONSPACE_LIMIT
D. LOG_ARCHIVE_TRACE and JAVA_MAX_SESSIONSPACE_SIZE

Answer: B

Question: 5
You use the following query to monitor the Shared Servers:

SQL>SELECT maximum_connections "MAX CONN",
servers_started "STARTED",
servers_terminated "TERMINATED",
servers_highwater "HIGHWATER" FROM v$shared_server_monitor;
MAX CONN STARTED TERMINATED HIGHWATER
--------------- ---------- ---------------- ----------
3 6 6 3

With every execution, you observe that the values of the STARTED and TERMINATED columns
are growing quickly. What action would you take?

A. The number of Shared Server sessions is too high, so you should decrease the value of
    SHARED_SERVER_SESSIONS.
B. The number of Shared Servers is too high, so you should decrease the number of Shared
    Servers.
C. No action is required because this growth is a normal behavior in a Shared Server
    environment.
D. The number of Shared Servers is too low, so you should increase the number of Shared
    Servers.

Answer: D

Question: 6
As part of the performance-tuning activity, you executed the utllockt.sql script. What does this
script do? (Choose all that apply.)

A. populates the v$locked_object view
B. creates the dba_locks and dba_blockers views
C. displays information about the sessions that are waiting for locks
D. displays information about the locks for which sessions are waiting

Answer: C, D

Question: 7
Which two statements are valid regarding row migration? (Choose two.)

A. It can be reduced by decreasing the PCTFREE value.
B. The row migration causes more than one block to be read when a row is accessed.
C. It can be resolved by using the ANALYZE TABLE...ESTIMATE STATISTICS command.
D. It is caused by inserting a row into an empty data block where the row is too large to fit into
    that empty block.
E. It is caused by an UPDATE statement that increases the data in a row so that the new larger
    version of the row no longer fits in its data block.

Answer: B, E

Question: 8

While tuning the shared server in your database, you executed the following command:

SQL> ALTER SYSTEM SET DISPATCHERS = '(PROTOCOL=TCP)(DISPATCHERS=5)';

How does this statement affect the shared server configuration? (Choose two.)

A. If there are currently less than five dispatchers, new dispatchers are created.
B. If there are currently more than five dispatchers, additional dispatchers are terminated after the connected users disconnect.
C. If there are currently more than five dispatchers, Oracle server terminates additional dispatchers immediately.
D. If there are currently less than five dispatchers, Oracle server terminates all of them and creates five new dispatchers.

Answer: A, B

Question: 9
Which two actions could you perform to reduce busy buffer waits caused on segment header blocks? (Choose two.)

A. configure multiple I/O slaves
B. increase the number of free lists
C. increase the size of the redo log buffer
D. use multiple database writer (DBWn) processes
E. use the Automatic Segment Space Management feature

Answer: B, E

Question: 10
In an online transaction processing (OLTP) system, which is configured in shared server mode, you execute the following query to examine the sort performance:

SQL> SELECT d.value "Disk", m.value "Mem", (d.value/m.value)*100 "Ratio"
FROM v$sysstat m, v$sysstat d
WHERE m.name='sorts (memory)' and d.name='sorts (disk)';
Disk Mem Ratio
---------- ---------- ---------
23 206 11.165049

What action would you take after observing the output?

A. increase the value of SORT_AREA_SIZE
B. increase the size of the temporary tablespace
C. increase the value of PGA_AGGREGATE_TARGET
D. no action required because the ratio is perfect for an OLTP system

Answer: A

Question: 11
You issued the following query on the V$INSTANCE_RECOVERY dynamic performance view:

SELECT target_mttr, estimated_mttr FROM v$instance_recovery;

You noticed that the value for the ESTIMATED_MTTR column is greater than the value for the TARGET_MTTR. What would happen in this scenario?

A. A checkpoint is initiated.
B. A warning is written in the alert log file.
C. The number of archiver processes is increased.
D. The control file is updated to mark another checkpoint record.
E. DBWn is requested to write enough dirty buffers to allow the ESTIMATED_MTTR to fall back below the TARGET_MTTR.

Answer: E

Question: 12
View the Exhibit and examine the execution plan of the following query:

SQL> SELECT s.prod_id
2 , t.fiscal_month_number
3 , sum(s.amount_sold) AS sum_amount
4 FROM sales s, times t
5 WHERE s.time_id = t.time_id
6 AND t.fiscal_year = 2000
7 GROUP BY s.prod_id, t.fiscal_month_number
8 ORDER BY s.prod_id, t.fiscal_month_number;

Which statement regarding the SALES_PROD object in the execution plan is true?
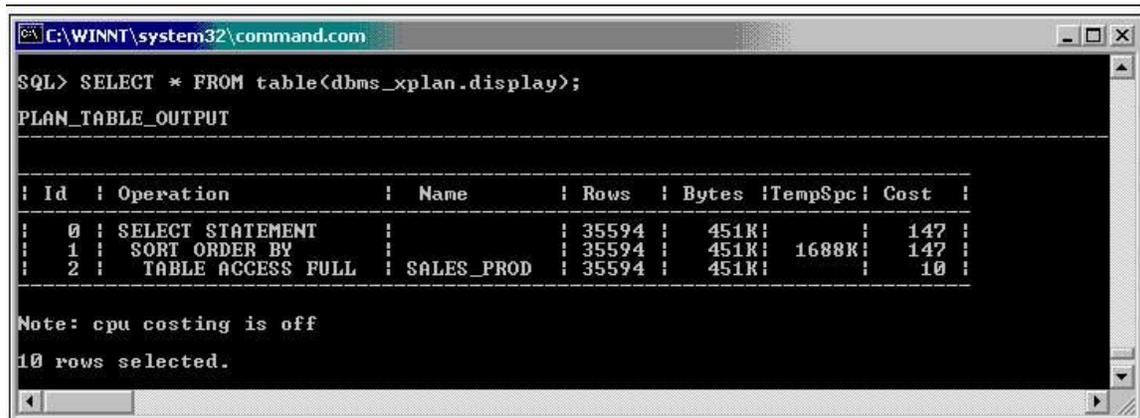
A. It is an index.
B. It is a complex view.
C. It is a mapping table.
D. It is a materialized view.
E. It is an index-organized table (IOT).

Answer: D

Question: 13
Which type of lock is held by SQL statements and PL/SQL objects in the Library Cache?

A. breakable parse locks
B. row-level (TX type) lock
C. table-level manual locks
D. table-level (TM type) locks

Answer: A

Question: 14
The steps involved in copying statistics from the source database to the target database are given below in random order:

1) Export the statistics table from the source database and then import it into the target database.
2) Copy the statistics to a table in the source database.
3) Create a table to hold the statistics in the source database.
4) Copy the statistics into the data dictionary in the target database.
5) Run the Oracle-supplied csminst.sql script to create a schema and a table to hold the statistics in the source database.
6) Collect the system statistics in the target database.

What is the correct order to accomplish this task?

A. 3, 2, 4, 1; 5 and 6 are not required
B. 5, 2, 1, 4; 3 and 6 are not required
C. 3, 2, 1, 4; 5 and 6 are not required
D. 5, 2, 1, 6; 3 and 4 are not required
E. 3, 4, 2, 1; 5 and 6 are not required

Answer: C

Question: 15
Which statement correctly describes the purpose of the REFRESH COMPLETE option when used with materialized view?

A. Truncates the materialized view and repopulates it with the data from the base table in the query
B. Populates the materialized view with the data from the base tables that has changed since the last refresh
C. Populates the materialized view with the data from the base tables after each commit to any of the base tables
D. Truncates the materialized view and populates it with the data from the base tables that has changed since the last refresh

Answer: A

Question: 16
View the Exhibit and examine the TKPROF output. While studying the TKPROF output, you find that the number of disk reads is close to the total number of buffers used (query + current) for some statements. What is the reason for this?

call	count	cpu	elapsed	disk	query	current	rows
Parse	1	0.00	0.00	0	0	0	0
Execute	1	0.00	0.00	0	0	0	0
Fetch	275	0.05	0.09	810	579	300	4096
total	277	0.05	0.09	810	579	300	4096

A. Shared pool is not large enough
B. Database buffer cache is not large enough

C. Too few DBWn processes
D. Fragmentation in the database buffer cache

Answer: B

Question: 17
You have the following requirements:

- Your transaction only queries the table, and requires a consistent set of the table's data for the
  duration of the transaction.
- You can hold up other transactions that try to update the locked table, until all transactions that
  hold SHARE locks on the table either commit or roll back.
- Other transactions may acquire concurrent SHARE table locks on the same table, also allowing
  them the option of transaction-level read consistency.

Which type of locking strategy would you select to achieve this?

A. SHARE
B. EXCLUSIVE
C. ROW SHARE
D. ROW EXCLUSIVE

Answer: A

Question: 18
You are working in a dedicated server environment. You discovered that your database
performance is degraded due to a huge number of sort operations that are using the disk
space instead of the memory are a. The WORKAREA_SIZE_POLICY parameter is set to AUTO.
What would be one of the solutions to improve database performance?

A. setting the *_AREA_SIZE parameters to higher values
B. setting the PGA_AGGREGATE_TARGET parameter to a higher value
C. setting the SORT_AREA_RETAINED_SIZE parameter to a lower value
D. setting the SORT_AREA_RETAINED_SIZE parameter to a higher value

Answer: B

Question: 19
You observed a lot of unused space below the high-water mark for a table. Which three methods
could you use to deallocate the unused space? (Choose three.)

A. follow the EXPORT, DROP, and IMPORT method
B. use the DBMS_SPACE.FREE_BLOCKS procedure
C. follow the EXPORT, TRUNCATE, and IMPORT method
D. use MOVE with the ALTER TABLE command to move the table
E. use DEALLOCATE UNUSED with the ALTER TABLE command

Answer: A, C, D

Question: 20
Which two statements are correct regarding users in a shared server environment? (Choose two.)

A. Connections can share dispatcher processes.
B. Multiplexing and pooling can be used for user connections.
C. It is advantageous for sessions performing batch operations.

D. The CPU overhead is reduced for each individual user request.
E. Multiple users cannot share the dispatcher processes to access the Oracle server.

Answer: A, B

Question: 21
Your database is running in Shared Server mode. You have set the LARGE_POOL_SIZE
parameter to 1 MB. While the database is functioning, you realize that the large pool is running
out of space because of the increasing number of user requests. What will the effect be if there is
no more free space left in the large pool to store new information?

A. The Shared Pool will be used to store the new information.
B. The large pool will be dynamically resized to store the new information.
C. User requests will fail but the database instance will continue to function.
D. The session will be switched to dedicated server mode and the PGA will be used.
E. The current information will be moved from the large pool to the Shared Pool, and the Shared
    Pool will continue to be used.

Answer: C

Question: 22
Smith is a DBA with XYZ Corp. There are 500 data entry users in an online transaction
processing (OLTP) environment. The current response time is 15-20 seconds. The company
wants the response time to be brought down to 10 seconds or less. What is the first thing that
Smith should do to diagnose the cause of the slow response time?

A. Determine whether there is contention for locks
B. Determine whether the file system is fast enough
C. Determine whether more memory needs to be allocated to PGA
D. Determine whether more memory needs to be allocated to SGA
E. Determine whether the slow response is wait bound or CPU bound

Answer: E

Question: 23
View the Exhibits and compare the baseline and current Load Profile statistics in the statspack
report. The customer had collected baseline statistics by using statspack six months
ago. The customer is facing performance issues and you have collected the current statistics.
Which option would you use on the basis of Load Profile statistics to improve the performance?

Baseline Statistics

Load Profile
~~~~~~~~~~~~~~~

```
                                Per Second        Per Transaction
                              ----------------    ----------------
              Redo size:         580,159.97            10,544.16
           Logical reads:        137,971.77             2,507.58
           Block changes:          5,516.11               100.25
           Physical reads:         8,969.62               199.37
           Physical writes:          768.78                13.97
```

Current Statistics

Load Profile
~~~~~~~~~~~~~~~

```
                                Per Second        Per Transaction
                              ----------------    ----------------
              Redo size:         202,029.61             3,787.10
           Logical reads:         86,928.46             1,629.50
           Block changes:          1,339.94                25.12
           Physical reads:        19,486.81               271.56
           Physical writes:          900.24                13.99
```

A. Enable query rewrite
B. Enable automatic PGA
C. Increase cursor sharing
D. Increase the size of buffer cache

Answer: D

Question: 24
You have many users complaining about slow inserts into a large table. While investigating the reason, you find that the number of indexes on the table is high. You want to find out which indexes are not being used. Which method would you follow to achieve this?

A. enable index monitoring and query the DBA_OBJECTS view
B. enable index monitoring and query the DBA_INDEXES view
C. enable index monitoring and query the V$OBJECT_USAGE view
D. enable index monitoring and query the DBA_INDEXTYPE_COMMENTS view

Answer: C

Question: 25
In the parameter file of your production database, the LOG_CHECKPOINT_INTERVAL parameter is set to 3600. When does the checkpoint occur due to this setting?

A. When the number of transactions exceeds the specified value
B. When the specified number of database blocks are written to the data files
C. When a specified amount of time (in seconds) is exceeded after the last checkpoint
D. When the number of dirty buffers in the database buffer cache exceeds the specified value
E. When the number of operating system blocks between an incremental checkpoint and the last
   block written to the redo log is equal to or greater than the specified value

Answer: E

Question: 26
Which three types of objects can you keep in the shared pool by using the
dbms_shared_pool.keep procedure? (Choose three.)

A. Types
B. Triggers
C. LOB segments
D. Sql cursor objects
E. PL/SQL procedures

Answer: B, D, E

Question: 27
After monitoring the Shared Server environment, you decided to increase the number of Shared
Servers available for use by the database by increasing the value of the initialization parameter
MAX_SHARED_SERVERS. Which condition requires this change in the parameter value?

A. The number of rows in the V$CIRCUIT view is equal to the value of the
   SHARED_SERVER_SESSIONS parameter.
B. The V$SHARED_SERVER_MONITOR view shows a high value for the SERVERS_STARTED
   column.
C. The V$QUEUE view shows an increasing number of requests in the queue and the value of
   WAIT column is increasing.
D. The V$SHARED_SERVER_MONITOR view shows that the value of the
   SERVERS_STARTED column is equal to the value of the MAX_SHARED_SERVERS
   parameter.

Answer: C

Question: 28
You executed the following query:

SQL> SELECT department_name,last_name
FROM employees E,departments D
WHERE E.department_id=D.department_id;
The EMPLOYEES and DEPARTMENTS tables have no indexes and the HASH_JOIN_ENABLED
parameter is set to false. From the following list of actions, select the correct sequence on the
basis of the order of their execution:

1)Sort each row sources separately
2)Perform full tablescans of EMPLOYEES and DEPARTMENTS tables
3)Merge the sorted sources

A. 2, 1, 3
B. 1, 2, 3
C. 2, 3, 1
D. 1, 2; 3 is not required

Answer: A

Question: 29
Oracle Shared Server is best used when _____ and _____. (Choose two.)

A. The database is experiencing database-intensive work
B. The database is primarily used for batch or DSS operations
C. CPU usage on your machine is consistently at 90% or higher
D. An OLTP application is running on a machine approaching memory resource limits
E. You do not need to scale up the number of concurrent connections to the database
F. The database is primarily used for an interactive application where dedicated servers are
   mainly idle

Answer: D, F

Question: 30
While querying the V$WAITSTAT data dictionary view, you find that the COUNT and TIME is very large for the class SEGMENT HEADER. How would you describe this problem?

A. Some segments are very large in size.
B. There is a contention in the least recently used (LRU) list.
C. The tablespace where the segments reside is very small.
D. The segment headers for some segments are very small.
E. There are not enough free list groups for certain segments.

Answer: E

Question: 31
Which two statements are valid regarding index clusters? (Choose two.)

A. Index clusters store null keys.
B. Cluster indexes are likely to be smaller than normal indexes on the same set of keys.
C. Index clusters are better than hash clusters when the number of key values is predictable.
D. Index clusters are better than hash clusters when queries use equality predicate on key
   column.

Answer: A, B

Question: 32
You find a DBA using the DBMS_STATS package to back up statistics before analyzing objects. Which two reasons could trigger this action? (Choose two.)

A. To restore old statistics
B. To conserve space in the database
C. To use the backed up statistics after a bulk load
D. To study changes in data characteristics over time
E. To use backed up statistics for a rule-based optimizer

Answer: A, D

Question: 33
You executed the following SQL command in your production database. The results of this command are displayed in the Exhibit:

SQL> SELECTshared_pool_size_for_estimate AS pool_size,
estd_lc_size,
estd_lc_time_saved
FROM v$shared_pool_Advice;
View the Exhibit.

Which statements are true regarding the result? (Choose all that apply.)

```
POOL_SIZE   ESTD_LC_SIZE ESTD_LC_TIME_SAVED
----------- ------------ ------------------
         24            3                 45
         32            3                 45
         40            3                 45
         48            3                 45
         56            3                 45
         64            3                 45
         72            3                 45
         80            3                 45
         88            3                 45
         96            3                 45
```

A. There are no additional hits on the size range for the library cache.
B. The estd_lc_size column value indicates the current size of the Shared Pool.
C. The pool_size column value indicates the range of the Shared Pool size that is estimated.
D. The estd_lc_time_saved column values are same for all pool sizes thereby indicating the
    Shared Pool size needs to be increased.

Answer: A, C

Question: 34
You set PGA_AGGREGATE_TARGET to a nonzero value. Which two statements are true
regarding this setting? (Choose two.)

A. The SQL workarea is automatically tuned by the instance.
B. The automatic SQL execution memory management is enabled.
C. The *_AREA_SIZE parameters must be set explicitly in the parameter file.
D. The maximum size of the Program Global Area (PGA) is equal to the value of
    PGA_AGGREGATE_TARGET parameter.
E. The SQL workarea must be tuned manually by using the *_WORK_AREA parameters.

Answer: A, B

Question: 35
Which two statements are true regarding the statistics used by the optimizer to generate an
optimal plan? (Choose two.)

A. The statistics collected through a histogram are not useful for better selectivity estimates of
    unevenly distributed data.
B. Dynamic sampling is preferred when the cost of collecting the sample statistics outweighs the
    advantage of a better execution plan.
C. The maximum and minimum values collected by the column statistics leads to an optimal plan
    if there is an even distribution of data between those two values.
D. The statistics collected through a histogram are useful if the data distribution is skewed
    between minimum and maximum values.

Answer: C, D

Question: 36

You created a mapping table on an index-organized table (IOT). What could be the reason for creating a mapping table?

A. You want to add the IOT to a cluster.
B. You want to keep track of the overflowing segments.
C. You want to store all the non-key columns of the IOT.
D. You want to create a bitmapped secondary index on the IOT.

Answer: D

Question: 37
Which statement is true regarding the TKPROF output when the AGGREGATE=Y option is used?

A. It aggregates the statistics for all recursive SQL statements.
B. It aggregates the statistics for all SQL statements in the trace file.
C. It aggregates the statistics for multiple uses of the same SQL text.
D. It aggregates the statistics for SQL statements for all users except SYS.

Answer: C

Question: 38
In your database the SORT_AREA_SIZE and HASH_AREA_SIZE are set to small values. Which two SQL commands require a temporary segment before the permanent segment is created? (Choose two.)

A. CREATE INDEX
B. CREATE UNDO TABLESPACE
C. CREATE ROLLBACK SEGMENT
D. ENABLE UNIQUE KEY CONSTRAINT

Answer: A, D

Question: 39
In your production environment:

- the database is running in ARCHIVELOG mode
- there are three online redo log groups
- there are two members in each group that are placed on different disks
- the archived redo log files are placed on disks other than the disks where the members of online
  redo log groups are placed

While performing transactions, users complain that transactions take a long time to complete. On investigation, you find that the STATUS column of the V$LOG data dictionary view shows the value ACTIVE, for two of the online redo log groups, thus causing the performance to be degraded. Which action would you perform to increase the performance?

A. Increase the size of the redo log buffer
B. Increase the number of online redo log groups
C. Increase the number of LogWriter (LGWR) processes
D. Increase the number of members in each online redo log group

Answer: B

Question: 40

View the Exhibit and compare the baseline and current data dictionary cache statistics in statspack reports. The DBA had collected baseline statistics by using statspack six months ago. For the purpose of performance tuning, the DBA collects the current statistics. The database system does not have a heavy insert load. What recommendation would you make on the basis of the statistics provided in the Exhibit?

A. Configure keep pool
B. Increase shared pool size
C. Configure recycle pool
D. Configure shared pool reserved area

Answer: B

Question: 41
While the DBWn process writes blocks from the database buffer cache to data files, you want the Oracle database to detect any kind of corruption that is caused by the underlying disk or storage system. Which task would you perform to achieve this?

A. ensure that the DB_BLOCK_CHECKING parameter is set to TRUE
B. ensure that the DB_BLOCK_CHECKSUM parameter is set to TRUE
C. ensure that the LOG_CHECKPOINT_TO_ALERT parameter is set to TRUE
D. use Recovery Manager (RMAN) Block Media Recovery to check the corruption automatically
E. ensure that the DB_BLOCK_CHECKING and DB_BLOCK_CHECKSUM parameters are both set to TRUE

Answer: B

Question: 42
You need to identify the top five events with the highest average wait time during the last 24 hours. Which option would you use to get the desired information?

A. run utlbstat/utlestat SQL scripts
B. create and use statspack report
C. query the V$SYSSTAT dynamic performance view
D. query the V$WAIT_EVENT dynamic performance view

Answer: B

Question: 43
Which two statements are correct guidelines regarding the process of choosing the number of dispatchers for the database? (Choose two.)

A. Kill the dispatcher for the user at the operating system (OS) level if the user causes one of the shared servers waiting.
B. Increase the number of dispatchers when the average wait time for users sessions is found to be steadily increasing.
C. When a large number of clients are connected through the same dispatcher, increase the number of dispatchers so that current users get service from new dispatchers.
D. The number of dispatchers for a protocol can be set irrespective of the value of the PROCESSES parameter.
E. Increase the number of dispatchers when the busy rate is over 50% and the number of clients for a dispatcher is high indicating high contention.

Answer: B, E

Question: 44
Which three statements correctly describe the causes of contention for Transaction Enqueue (TX)? (Choose three.)

A. Waits due to create view or alter view operations
B. Waits due to Unique or Primary Key constraint enforcement
C. Waits due to rows being covered by the same Bitmap index fragment
D. Waits due to insufficient interested transaction list (ITL) slots in the block
E. Waits due to executing ANALYZE TABLE..COMPUTE STATISTICS or VALIDATE STRUCTURE command

Answer: B, C, D

Question: 45
You are working in an OLTP environment. You wish to reduce the virtual memory requirements for the Oracle instance due to Operating system paging or swap overheads. You need to determine whether the buffer cache is oversized and could tolerate being reduced. In which two cases can you afford to reduce the buffer cache size? (Choose two.)

A. If the cache hit ratio is very high
B. If there are no waits for free buffers
C. If the rate of physical I/O is very high
D. If the number of full table scans is very low
E. If the number of index accesses is very high

Answer: A, B

Question: 46
Which two statements correctly describe the performance benefits of having large block size in a decision support system (DSS) environment? (Choose two.)

A. Large blocks reduce block contention.
B. Large blocks are suitable for random access.
C. Large blocks are suitable for storing very large rows.
D. Large blocks lead to fewer I/Os when traversing index branches.

Answer: C, D

Question: 47
You are working in an online transaction processing (OLTP) environment. The NORTH_SALES table contains one million rows. In which case would you issue the following command:

SQL> ALTER TABLE north_sales STORAGE(BUFFER_POOL recycle);

A. The keep pool is not configured.
B. The automatic PGA is not configured.
C. The NORTH_SALES table is accessed once a week.
D. The CURSOR_SPACE_FOR_TIME initialization parameter is set to FALSE.
E. All the queries on the NORTH_SALES table are rewritten using a materialized view.

Answer: C

Question: 48
You want to condense the space of an index on one of the important tables in your database. You used the ALTER INDEX COALESCE statement to achieve this task. Which requirement would have forced you to use this method instead of rebuilding the index? (Choose all that apply.)

A. You have a disk space constraint.
B. You want to create a new tree for the index.
C. You want to move the index to another tablespace.
D. You want to free up fragmented index leaf blocks quickly, for use.

Answer: A, D

Question: 49
View the Exhibit and examine the command used to create the ZONEDATA table.
The table contains a million rows for zonewise analysis in DSS system. You decide to prepare an index on zone column to enhance the performance of the queries on ZONE column. Which type of index would you select?

```
CREATE TABLE zonedata (col_id NUMBER(5) NOT NULL,
  col_dt date,
  zone VARCHAR2(5) CONSTRAINT z_check CHECK(zone IN('EAST','WEST','NORTH','SOUTH')),
  amount NUMBER(12,2) CONSTRAINT amt_chk CHECK(amount > 0),
  cust_id VARCHAR2(5) NOT NULL,
  category VARCHAR2(10) )
```

A. Bitmap index
B. Reverse key index

C. Compressed index
D. Normal B-Tree index

Answer: A

Question: 50
You create a stored outline and apply it for use. What strategy would you adopt to verify that the stored outline is being used for a statement?

A. verify OL$,OL$HINTS, and OL$NODES tables after running the statement with
    USE_STORED_OUTLINES set
B. examine the V$SQL dynamic performance view after running the statement with
    theUSE_STORED_OUTLINES set
C. compare the explain plan output for the statement when running with and without the
    USE_STORED_OUTLINES set
D. compare the statspack output by taking snaps before and after running the statement with the
    USE_STORED_OUTLINES set

Answer: C

Question: 51
You have the following outcome from the V$SYSSTAT view:

SQL> SELECT name, value FROM v$sysstat WHERE name = 'table fetch continued row';
NAME VALUE
-- ----------
table fetch continued row 308

What action would you take to have control of this value in the future?

A. Increase the NEXT setting for the tables with row migration problems
B. Increase the PCTUSED setting on the tables with row migration problems
C. Increase the PCTFREE setting on the tables with row migration problems
D. Increase the MAXEXTENTS setting on the tables with row migration problems

Answer: C

Question: 52
You have reserved memory within the Shared Pool by using the
SHARED_POOL_RESERVED_SIZE parameter to accommodate large space allocations while
compiling PL/SQL blocks and database triggers. While observing the performance of the
reserved area, you find that the value of the REQUEST_FAILURES column in the
V$SHARED_POOL_RESERVED view is continuously increasing. Which statement is true in this
scenario?

A. The parsing of the statement fails due to invalidations.
B. The execution of the statement fails due to invalidations.
C. The space allocated for the reserved area is not enough.
D. The syntax in the PL/SQL blocks is causing the compilation failure.

Answer: C

Question: 53
You executed the following query on the V$SYSTEM_EVENT view:

SQL> SELECT event, total_waits, time_waited FROM v$system_event; The result of this query displays the total waits and time waited for events as shown in the Exhibit. View the Exhibit. The query provides the total waits and time waited for the events of _____.

```
EVENT                                    TOTAL_WAITS TIME_WAITED
---------------------------------------- ----------- -----------
latch free                                       135           0
pmon timer                                     88125    25179657
rdbms ipc reply                                   22         116
rdbms ipc message                             547975   119952883
enqueue                                           32           4
async disk IO                                     47          71
control file sequential read                 1176869       60567
control file parallel write                    81917        9905
refresh controlfile command                        1           1
control file heartbeat                             1         411
checkpoint completed                               1           7
buffer busy waits                             142793    14726228
recovery read                                      2           0
log file sequential read                          12          49
log file single write                              7           0
log file parallel write                        31127        1247
LGWR wait for redo copy                          806           6
log file switch (archiving needed)            143069    14752404
log file switch completion                         1          36
log file sync                                      4           0
db file sequential read                          928         661
db file scattered read                           112         147
db file single write                              11           0
db file parallel write                          4582         140
direct path read                                  44           0
direct path write                                 20           0
instance state change                              1           0
```

A. The user session
B. The current instance
C. The currently running queries
D. The currently connected users
E. All the instances of the database

Answer: B

Question: 54
The buffer cache advisory feature is off in your database. You execute the following command to enable the Buffer Cache advisory feature in the database instance:

SQL> alter system set DB_CACHE_ADVICE = on;
After executing the command, you receive the following error:
ORA-04031: unable to allocate shared memory.

What could be the reason for this error?

A. The database instance is in OPEN mode.
B. The Shared Pool size is not enough to turn on the advisory feature.

C. A large number of users are connected, thereby causing memory overhead.
D. The Database Buffer Cache size is not enough to turn on the advisory feature.

Answer: B

Question: 55
Which three techniques should a DBA use to reduce the overhead on library cache? (Choose three.)

A. Flush shared pool at regular intervals
B. Configure Keep and Recycle buffer pools
C. Ensure that the shared pool is sufficiently sized
D. Set the CURSOR_SHARING parameter to SIMILAR or FORCE
E. Set CURSOR_SPACE_FOR_TIME initialization parameter to FALSE
F. Avoid data definition language (DDL) statements that cause invalidations
G. Use literals in SQL statements when the CURSOR_SHARING parameter is set to EXACT

Answer: C, D, F

Question: 56
You executed the following query to know about users and their consumer groups:

SQL> SELECT * FROM DBA_RSRC_CONSUMER_GROUP_PRIVS;
GRANTEE GRANTED_GROUP GRANT_OPTION INITIAL_GROUP
--------------- --------------------------- -------------------- --------------------
SCOTT LOW_GROUP NO NO
PUBLIC LOW_GROUP NO NO
PUBLIC DEFAULT_CONSUMER_GROUP YES YES
SYSTEM SYS_GROUP NO YES

If the user SCOTT starts a session, to which consumer group will the session belong?

A. The session belongs to LOW_GROUP.
B. The session belongs to SYS_GROUP.
C. The session belongs to DEFAULT_CONSUMER_GROUP.
D. The session does not belong to any group because the user is not assigned to any group.

Answer: C

Question: 57
You executed the following command to observe the index statistics:

SQL> SELECT name,height,blocks,lf_rows,del_lf_rows,lf_blks,distinct_keys
2 FROM index_stats
3 WHERE name='CUIND';
NAME HEIGHT BLOCKS LF_ROWS DEL_LF_ROWS LF_BLKS DISTINCT_KEYS
------- ---------- ---------- ------------ --------------------- ---------- ---------------------
CUIND 2 24 8192 4096 18 32

Which statement is true regarding the CUIND index?

A. The index is very rarely used.
B. The index is a candidate for rebuilding.
C. The index has most of the leaf blocks packed with rows.
D. Parts of the index are accessed more frequently than others.

Answer: B

Question: 58
In your database, Smith and John work on the HR schem a. They connect to the HR schema and update the dept table simultaneously from their respective sessions:

Smith's session:
SQL> UPDATE dept SET deptno=60 WHERE loc='NEW YORK';
1 row updated.
John's session:
SQL> UPDATE dept SET deptno=80 WHERE loc='BOSTON';

In Smith's session, the UPDATE statement is successful whereas in John's session the UPDATE statement is waiting for Smith to complete his transaction. What could be the reason for this?

A. The value of the DML_LOCKS parameter is changed.
B. It is due to the default locking mechanism in the Oracle database.
C. Both of them logged in as HR and the HR user has DBA privileges.
D. The value of the ROW_LOCKING parameter is changed to intent.
E. Both of them logged in as HR and the HR user is not granted the RESOURCE role.

Answer: D

Question: 59
You receive complaints from users regarding the high waiting time for their transactions. On investigation, you find that some users are not committing their transactions even though their sessions do not perform any activity for a long time. Which strategy would prevent this locking problem?

A. Decrease the IDLE_TIME limit in the profile of the blocking users
B. Use the resource manager to control the idle time for blocking users
C. Instruct the blocking users to manually lock the target table in EXCLUSIVE mode
D. Set the limit in the profile of blocking users to control the number of blocks to be accessed in a session

Answer: A

Question: 60
A customer has various dictionary-managed tablespaces with past dat a. The extents in dictionary-managed tablespaces are not uniformly sized. Newly created tablespaces are locally managed. You want to make all the tablespaces locally managed. What approach would you use to change the dictionary-managed tablespaces to locally managed tablespaces to get all the performance benefits of local extent management?

A. use only the DBMS_SPACE_ADMIN.TABLESPACE_MIGRATE_TO_LOCAL procedure
B. create new locally managed tablespaces and move data from existing dictionary-managed tablespaces to locally managed tablespaces
C. use the DBMS_SPACE_ADMIN.TABLESPACE_MIGRATE_TO_LOCAL procedure, followed by DBMS_SPACE_ADMIN.TABLESPACE_FIX_SEGMENT_STATES
D. use the ALTER TABLESPACE command to change extent management to local, and then use the DBMS_SPACE_ADMIN.TABLESPACE_MIGRATE_TO_LOCAL procedure

Answer: B

Question: 61
View the Exhibit and examine the alert log file. The Service Level Agreement for the production database includes that the database down-time should not be more than 10 minutes. What performance-related issue could you identify by examining the alert log file?

```
ALTER DATABASE OPEN\par
Thu Dec 29 12:14:44 2005\par
Beginning crash recovery of 1 threads\par
Thu Dec 29 12:14:44 2005\par
Started first pass scan\par
Thu Dec 29 12:14:45 2005\par
Completed first pass scan\par
 3491 redo blocks read, 69 data blocks need recovery\par
Thu Dec 29 12:14:45 2005\par
Started recovery at\par
 Thread 1: logseq 136, block 89818, scn 0.0\par
Recovery of Online Redo Log: Thread 1 Group 3 Seq 136 Reading mem 0\par
  Mem# 0 errs 0: D:\\INSTALLATIONS\\ORA9I\\ORADATA\\ORA9I\\REDO03.LOG\par
Thu Dec 29 12:14:45 2005\par
Ended recovery at\par
 Thread 1: logseq 136, block 93309, scn 0.36385921\par
 69 data blocks read, 69 data blocks written, 3491 redo blocks read\par
Crash recovery completed successfully\par
Thu Dec 29 12:44:46 2005\par
Thread 1 advanced to log sequence 137\par
Thread 1 opened at log sequence 137\par
  Current log# 1 seq# 137 mem# 0: D:\\INSTALLATIONS\\ORA9I\\ORADATA\\ORA9I\\REDO01.LOG\par
Successful open of redo thread 1.\par
Thu Dec 29 12:44:46 2005\par
SMON: enabling cache recovery\par
Thu Dec 29 12:44:47 2005\par
Undo Segment 1 Onlined\par
Undo Segment 2 Onlined\par
Undo Segment 3 Onlined\par
Undo Segment 4 Onlined\par
Undo Segment 5 Onlined\par
Undo Segment 6 Onlined\par
Undo Segment 7 Onlined\par
Undo Segment 8 Onlined\par
Undo Segment 9 Onlined\par
Undo Segment 10 Onlined\par
Successfully onlined Undo Tablespace 1.\par
Thu Dec 29 12:44:47 2005\par
SMON: enabling tx recovery\par
Thu Dec 29 12:44:50 2005\par
Database Characterset is WE8MSWIN1252\par
replication_dependency_tracking turned off (no async multimaster replication found)\par
Completed: ALTER DATABASE OPEN
```

A. High log switch frequency
B. High checkpoint frequency
C. 69 data blocks are corrupted
D. Excessive crash recovery time

Answer: D

Question: 62
Examine the following query:

SQL> SELECT num_rows, blocks, empty_blocks as empty, avg_space
2 FROM dba_tables
3 WHERE owner='HR'
4 AND table_name='DEPARTMENTS';
NUM_ROWS BLOCKS EMPTY AVG_SPACE
---------------- ------------ ---------- -------------
12145 175 13 1250

What does the number 13 displayed in the EMPTY column indicate?

A. The number of blocks below the high water mark of the table

B. The number of empty blocks that have not yet been used in the table
C. The number of empty blocks containing data that have been deleted from the table
D. The number of blocks above the high water mark and the number of empty blocks containing data that have been deleted from the table

Answer: B

Question: 63
ABC Corp. plans to build an online transaction processing (OLTP) system. The company plans to tune all the components during the development life cycle. From the list given below, identify the order in which the components should be tuned.

1. logical and physical model
2. cursor sharing
3. Oracle instance memory structures
4. operating system
5 data file I/O

A. 1, 2, 3, 5, 4
B. 1, 3, 2, 4, 5
C. 1, 2, 3, 4, 5
D. 1, 2, 5, 3, 4

Answer: A

Question: 64
Consider the following scenario:

- You are working on a production database; you cannot shut down and restart the instance frequently.
- During daytime, you need to maintain a higher cache hit ratio.
- During nighttime, there are no online transactions and memory requirements are mainly for Oracle backup and restore operations.

Which strategy would you recommend to optimally use the memory?

A. Change SGA_MAX_SIZE by using the ALTER SYSTEM command during daytime and nighttime
B. Set the USE_INDIRECT_DATA_BUFFERS and CURSOR_SPACE_FOR_TIME initialization parameters to suit the buffer cache and large pool requirements in the respective time windows
C. Use the dynamic SGA feature; shrink and extend buffer cache and large pool as required, using the ALTER SYSTEM command during daytime and nighttime
D. Set the size of the buffer cache to the maximum size requirement and set the size of the large pool to the minimum size requirement

Answer: C

Question: 65
In your test database, you want to automatically collect database statistics at regular intervals. You execute the spauto.sql script to automate this task. What must you ensure before executing this script?

A. The LOCK_SGA parameter is set to TRUE.
B. The THREAD parameter is set to a nonzero value.

C. The OPTIMIZER_MODE parameter is set to CHOOSE.
D. The JOB_QUEUE_PROCESSES parameter is set to a nonzero value.
E. The AQ_TM_PROCESSES parameter is set to a value higher than one.

Answer: D

Question: 66
The database performance is degrading. You need to isolate the SQL statements that are contributing to the disk I/O load. Which diagnostic step would you recommend in the given scenario?

A. querying the V$SYSSTAT dynamic performance view to identify the SQL statements causing maximum I/O
B. query the V$SQL dynamic performance view to check the value for the EXECUTIONS column
C. using the findings of the Reorg Wizard to identify the SQL statements that are causing maximum I/O
D. using Oracle SQL Analyze application of the Oracle Tuning pack to identify the SQL statements that consume the most resources

Answer: D

Question: 67
Because of fragmentation, a frequently required large PL/SQL package is being flushed out of the library cache. Which strategy would you recommend to ensure the availability of the package in the library cache?

A. pin the package in your Private SQL Area
B. set the LOCK_SGA parameter within the package to ensure memory space availability for the package
C. load the package in the library cache during instance startup using a startup trigger and use the dbms_shared_ pool.keep procedure to keep the object in the shared pool
D. load the package in the library cache minimum 2 minutes after the instance startup, use the dbms_shared_ pool.keep procedure to keep the object in the shared pool, and ensure that you do not flush the shared pool

Answer: C

Question: 68
How would you identify the buffer cache blocks being used by an object?

A. by running a statspack report and using that report
B. by querying the DBA_OBJECTS data dictionary view
C. by querying the V$CACHE dynamic performance view
D. by querying the V$OBJECT_USAGE dynamic performance view

Answer: C

Question: 69
You noticed that the LGWR process is slow in freeing the redo log buffers. On querying the V$SYSTEM_EVENT view, you find that some time is spent on waiting for the log file switch (archiving needed) event. The archiving directories are not full. What should you do to minimize the waits? (Choose all that apply.)

A. Add redo log groups
B. Increase the number of database writer processes

C. Increase the number of archive log file destinations
D. increase the value for the LOG_BUFFER initialization parameter
E. increase the maximum number of archive processes by using the
   LOG_ARCHIVE_MAX_PROCESSES parameter

Answer: A, E

Question: 70
While doing a routine performance check of your test database, you find that a large number of statements are getting invalidated in the library cache, thereby causing the statements to be reparsed. What could be the reason for the invalidations?

A. Stored outlines are being used.
B. There is not enough free space in the Shared Pool.
C. The Shared Pool Reserved Space is not configured.
D. The structures of dependent objects have been modified.
E. Statistics have not been collected on the referenced columns.

Answer: D

Question: 71
Application developers are complaining about the poor response time of SQL statements on certain tables. On investigation, you find that the generation of recursive calls (in addition to SQL statements) to perform space management operations is the cause for the poor response time. Which of the solutions provided would solve the problem?

A. Using a bigger undo tablespace
B. Altering the database block size
C. Allocating extents to the tables to avoid dynamic extent allocation
D. Encouraging the use of anonymous PL/SQL blocks and executing SQL statements within them

Answer: C

Question: 72
Which two factors should be ensured by a DBA or an application developer while designing locking procedures to prevent locking problems? (Choose two.)

A. Explicit specification of index locks
B. Optimum level of data concurrency
C. Explicit specification of row-level locks
D. Locks held for the shortest possible time
E. Changes should be automatically committed by the system after an interval defined by the
   DBA

Answer: B, D

Question: 73
Your production database is running in ARCHIVELOG mode but the Archiver (ARCn) process is not started yet. What is the effect of this?

A. The database continues to function normally without any archived redo log files being created.
B. The Log Writer (LGWR) process archives the online redo log files in the absence of the ARCn
   process.
C. The ARCn process is automatically started when the LGWR process attempts to overwrite the
   online redo log files.

D. When an attempt is made to overwrite the unarchived online redo log, then LGWR will wait and all transactions will get stuck.

Answer: D

Question: 74
There are three users working on the same schema in your database. The techdt table in this schema does not have any indexes on it. All the three users execute UPDATE statements on the techdt table to update different rows. When the third user executes an UPDATE statement, it waits and does not update the table. You execute the following query to find the reason for this:

SQL> SELECT sid,type,id1,id2,lmode,request FROM v$lock WHERE type='TX';
SID TYPE ID1 ID2 LMODE REQUEST
---------- -------- ---------- ---------- ------------- -------------
11 TX 262150 45564 6 0
16 TX 393253 45294 6 0
18 TX 393253 45294 0 4

Why does the UPDATE statement of the third user wait?

A. The INITRANS parameter is set to a lower value.
B. The PCTFREE parameter is set to a higher value.
C. All the transaction slots up to MAXTRANS in a block are in use.
D. This is because the ROW_LOCKING parameter is set to INTENT.

Answer: C

Question: 75
In your database, Transaction 1 and Transaction 2 are waiting for rows locked by each other. This causes a deadlock. What would the Oracle server do in this situation? (Choose all that apply.)

A. Generate a trace file
B. Report an error in the alert log file
C. Terminate all the sessions that caused the deadlock
D. Rollback the statement in one of the sessions involved in the deadlock
E. Wait for three seconds and then roll back the statements that caused the deadlock to be detected

Answer: A, B, D

Question: 76
When a row in a table is chained or migrated, performance associated with that row degrades. What are the two reasons for this? (Choose two.)

A. Because the ROWID for migrated rows have changed
B. Because all the existing indexes on the table have become invalid
C. Because Oracle must scan more than one data block to retrieve the information for migrated rows
D. Because additional I/Os are performed when SELECT statements use an index to select migrated or chained rows

Answer: C, D

Question: 77

You want to compact the space of an index on one of the important tables in your database. You used the ALTER INDEX REBUILD ONLINE to achieve this task. Which requirement(s) would have forced you to use this method instead of coalescing the index? (Choose all that apply.)

A. You have a disk space constraint.
B. You want to create a new tree for the index.
C. You want to move the index to another tablespace.
D. You do not want concurrent operations on the base table.

Answer: B, C

Question: 78
View the Exhibit and examine the PGA advisory from the statspack report. You notice from the statspack report that PGA_AGGREGATE_TARGET is set too low for the current workload. The current value of PGA_AGGREGATE_TARGET is 24 MB. What is the optimal value for the PGA_AGGREGATE_TARGET parameter?

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
12	0.5	42.9	30.9	58.0	9
18	0.8	42.9	30.9	58.0	9
24	1.0	42.9	7.7	85.0	8
29	1.2	42.9	6.9	86.0	7
34	1.4	42.9	6.9	86.0	5
38	1.6	42.9	6.9	86.0	3
43	1.8	42.9	6.9	86.0	0
48	2.0	42.9	6.9	86.0	0
72	3.0	42.9	0.0	100.0	0
96	4.0	42.9	0.0	100.0	0
144	6.0	42.9	0.0	100.0	0
192	8.0	42.9	0.0	100.0	0

A. 72 MB
B. 12 MB
C. 29 MB
D. 192 MB

Answer: A

Question: 79
The values for the resource plan directives are set as follows:

SWITCH_GROUP = OLAP
SWITCH_TIME = 1000
SWITCH_ESTIMATE = True

What are the two implications of these values on a running session belonging to the consumer group for which these directives are set? (Choose two.)

A. These settings are useful to limit the resources consumed by short-running operations.

B. The resource manager switches the session to the OLAP group if the session is active for more than 1,000 seconds and after the operation the session becomes part of the OLAP group.
C. The resource manager can switch the session to the OLAP group before an operation even starts running.
D. The resource manager switches the session to the online analytical processing (OLAP) group if the session is active for more than 1,000 seconds and after the operation the session is moved back to the original group.

Answer: C, D

Question: 80
You created an index on the pdcost column of the PDLIST table in your database, and then executed the following command to find the total number of rows sorted until now:

SQL> SELECT name, value FROM v$sysstat WHERE name='sorts (rows)';
NAME VALUE
-- ----------
sorts (rows) 54762

Execution of which two queries would increase the sorts(rows) value? (Choose two.)

A. SELECT sum(pdcost) FROM pdlist;
B. SELECT max(pdcost) FROM pdlist;
C. SELECT DISTINCT pdcost FROM pdlist;
D. SELECT * FROM pdlist ORDER BY pdcost;

Answer: C, D

Question: 81
You noticed that the CUSTOMERS table is facing excessive row migration, thereby causing contention for free lists. What would you do to reduce row migration? (Choose all that apply.)

A. Increase the value for PCTFREE
B. Increase the value for PCTUSED
C. Decrease the value for PCTFREE
D. Reorganize the CUSTOMERS table
E. Increase the number of transaction slots

Answer: A, D

Question: 82
Users complain of slow response time in your online transaction processing (OLTP) system. You compare current statistics with baseline statistics. You notice that in the current statistics library cache, the hit ratio has decreased and there are request misses for shared pool reserved are a. Given below are the sizes of shared pool and shared pool reserved area:

SHARED_POOL_SIZE = 500 MB
SHARED_POOL_RESERVED_SIZE = 50 MB

What kind of strategy would you suggest for performance tuning?

A. Increase the size of only the shared pool
B. Decrease the size of the shared pool reserved area
C. Increase the size of only the shared pool reserved area

D. Increase the sizes of both the shared pool and the shared pool reserved area
E. Increase the size of the shared pool and decrease the size of the shared pool reserved area

Answer: D

Question: 83
Why does performance degrade when many UPDATE, INSERT, or DELETE statements are issued on a table that has an associated Bitmap index?

A. Some DML operations re-create the Bitmap index blocks.
B. The Bitmap index is rebuilt automatically after a DML operation.
C. The smallest amount of a bitmap that can be locked is a bitmap segment.
D. Additional time is taken to remove NULL values from the Bitmap index after a DML operation.

Answer: C

Question: 84
You executed the following command:

SQL> SELECT paddr,type,wait,totalq FROM v$queue;
PADDR TYPE WAIT TOTALQ
---------- ------------------ ---------- ----------
00 COMMON 852 6819
709E687C DISPATCHER 21 3444
709E6FFC DISPATCHER 12 3381

For the process 00, the wait column shows a significant increase in the value. Identify two possible solutions for this. (Choose two.)

A. Implement connection pooling
B. Increase the number of dispatchers
C. Increase the value of the MAX_SHARED_SERVERS parameter
D. Check the system memory capacity; a low system memory can cause the Shared Servers to run out of memory

Answer: C, D

Question: 85
The KEEP pool is full. You issued a query requiring buffers from the KEEP pool. What will happen?

A. The DEFAULT pool will be used instead.
B. The RECYCLE pool will be used instead.
C. The query will return an out-of-memory error.
D. Some of the existing object blocks will be aged out of the KEEP pool.

Answer: D

Question: 86
Which three actions on the TEST_TBS tablespace would result in tablespace checkpoint? (Choose three.)

A. ALTER TABLESPACE test_tbs READ ONLY;
B. ALTER TABLESPACE test_tbs END BACKUP;
C. ALTER TABLESPACE test_tbs BEGIN BACKUP;

D. ALTER TABLESPACE test_tbs OFFLINE NORMAL;
E. ALTER TABLESPACE test_tbs OFFLINE IMMEDIATE;

Answer: A, C, D

Question: 87
What would you do to reduce fragmentation within extents?

A. Avoid using MINEXTENTS clause while creating a segment
B. Reorganize the segments in the uniform extent size tablespaces
C. Defragment nonuniform extent size tablespaces by using export/import
D. Use locally managed tablespaces instead of dictionary-managed tablespaces

Answer: C, D

Question: 88
View the Exhibit.
In your database, you find that the fy05q1 table in a dictionary managed tablespace has excess space. Further, you find that the table consists of three extents. The size of the first extent is 100KB, the second is 200KB, and the third is 300K. The high water mark is in the middle of the second extent, and there is 400KB of unused space. You want to release the unused space so that it can be used by other segments. You execute the following command to achieve this task:

SQL> ALTER TABLE fy05q1 DEALLOCATE UNUSED;

What would this command do?
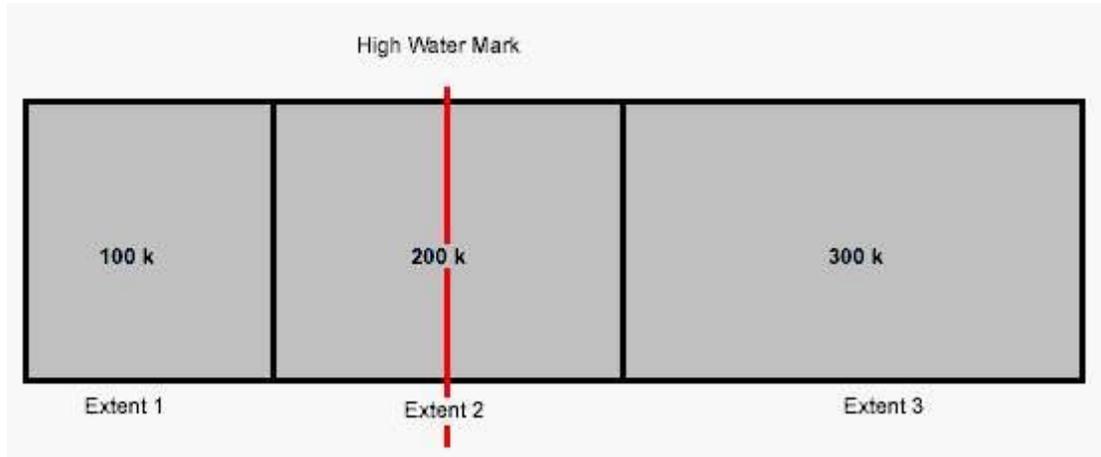
A. Deallocates the third and second extents
B. Deallocates the third extent; the second extent remains as it is
C. Deallocates the third extent, and the second extent is sized to 100KB
D. Returns an error because the unused space is above the high water mark

Answer: C

Question: 89
In your test database you have two applications (Application A and Application B) returning the same result set.

- Application A has buffer cache hit ratio of 99%, requires 100,000 logical reads, and 1,000 physical reads
- Application B has buffer cache hit ratio of 60%, requires 1,000 logical reads, and 400 physical reads

Which statement is correct?

A. No conclusion can be drawn without checking the response times of the applications.
B. Application B is better tuned than Application A, because the number of logical and physical reads in Application B is much less than that of Application A.
C. Application A is better tuned than Application B, because the ratio of physical reads to logical reads in Application A is lower than that of Application B.
D. Application A is better tuned than Application B, because the cache hit ratio in Application A is higher than that of Application B.

Answer: B

Question: 90
In the parameter file, you set the DB_CACHE_ADVICE parameter to READY. What is the effect of the parameter when starting the database instance?

A. The advisory is turned off and the memory allocated for the advisory is released.
B. The advisory is turned off and the memory is allocated for the advisory.
C. The advisory is turned on but the Buffer Cache advisory information is not collected.
D. The advisory is turned on and the Buffer Cache advisory information is collected in the V$DB_CACHE_ADVICE view.
E. The setting would result in an ORA-04031 error because this parameter must be set by using the ALTER SYSTEM command.

Answer: B

Question: 91
You need to size the KEEP pool. Which calculation would you use to determine the initial size of the KEEP pool?

A. 50% of the DEFAULT pool
B. 50% of the RECYCLE pool
C. the total number of blocks of all the candidate objects that you need in the KEEP pool / the number of objects that you need in the KEEP pool
D. total number of blocks of all the candidate objects that you need in the KEEP pool + additional space to allow for object growth

Answer: D

Question: 92
View the Exhibit and examine the parameters. Why would the DBA specify different values for the SESSIONS and SHARED_SERVER_SESSIONS parameters?

```
SQL> SHOW PARAMETER SESSION

NAME                               TYPE          VALUE
---------------------------------- -----------   -------------------------------
java_max_sessionspace_size         integer       0
java_soft_sessionspace_limit       integer       0
license_max_sessions               integer       0
license_sessions_warning           integer       0
logmnr_max_persistent_sessions     integer       1
mts_sessions                       integer       150
session_cached_cursors             integer       0
session_max_open_files             integer       10
sessions                           integer       170
shared_server_sessions             integer       150
```

A. To reserve sessions for the dispatcher processes
B. To reserve user sessions for dedicated connections
C. To reserve sessions for Oracle server's internal sessions
D. To reserve sessions for use when the database instance is hung

Answer: B

Question: 93
You decide to create an index-organized table in your database. The table would hold a large
number of rows and different departments would work only with the relevant rows in
the table concurrently. Which statement regarding the above table can enhance performance?

A. It can be partitioned.
B. It can be added to a cluster.
C. It can be created without a primary key.
D. It can be stored using a bitmap structure.

Answer: A

Question: 94
You are working on the performance tuning of your database. Which two types of performance
tuning-related information would an alert log file provide? (Choose two.)

A. Mean time to recover
B. Latch contention information
C. Instance recovery start and complete times
D. Sql statements that are consuming maximum resources
E. The top five sessions that are consuming maximum resources

Answer: A, C

Question: 95
Which three statements are true regarding STATSPACK? (Choose three.)

A. STATSPACK data should be separated from production data by putting it in its own
   tablespace.
B. Obsolete statistical data collected by STATSPACK cannot be purged. DBAs will need to drop
   the user PERFSTAT and reinstall STATSPACK.
C. To make performance comparisons from one day, week, or year to the next, there must be
   multiple snapshots taken over a period of time.

D. You can capture both the complete text of a high-load SQL statement and information on any SQL plan(s) associated with that statement in a STATSPACK report.
E. When a snapshot is taken, STATSPACK will automatically generate a report that displays the statistical output, including load profile, efficiency percentage of the instance, and top five wait events.

Answer: A, C, D

Question: 96
You increased the value of the OPTIMIZER_MAX_PERMUTATIONS parameter from 4000 to 8000. What would be the impact of this increment on the queries that are being executed? (Choose all that apply.)

A. The query parse time increases.
B. The query execution time increases.
C. The scope of the optimizer increases to select a good execution plan.
D. The scope of the optimizer increases to select between rule-based and cost-based optimization.

Answer: A, C

Question: 97
You have a trace file with you for analysis. How can you analyze the trace file and get most resource-intensive statements, aggregation of statistics, and inclusion or exclusion details of recursive calls?

A. By using the TKPROF utility
B. By using the ANALYZE command
C. By using the DBMS_TRACE package
D. By using the EXPLAIN PLAN command

Answer: A

Question: 98
View the Exhibit and examine the value set for initialization parameters. The customer had collected baseline statistics by using statspack six months ago. The current statistics indicate that workarea multipass executions have gone up substantially when compared with baseline statistics. Which recommendation would you make?

```
query_rewrite_enabled=FALSE
java_pool_size=33554432
large_pool_size=8388608
shared_pool_size=50331648
pga_aggregate_target=25165824
sort_area_size=524288
```

A. Enable query rewrite
B. Increase large pool size
C. Increase the size of shared pool
D. Increase bitmap merge area size
E. Increase the value of the PGA_AGGREGATE_TARGET initialization parameter

Answer: E

Question: 99
You execute the following query to diagnose the performance of the shared servers:

SQL> SELECT name "NAME", paddr, requests,
(busy / (busy+idle)) * 100 "%TIME BUSY", status
FROM V$SHARED_SERVER;
NAME PADDR REQUESTS %TIME BUSY STATUS
------- ---------- ---------------- ---------------- ----------------
S000 709E64BC 8077 .018732129 WAIT(ENQ)
S001 709E73BC 2953 .009255781 WAIT(COMMON)
S002 709E777C 1079 .158928369 WAIT(RECEIVE)
S003 709E7B3C 7617 .022872827 WAIT(COMMON)

What can you deduce from this result?

A. The user assigned to the shared server S000 is waiting for a lock resource.
B. The user assigned to the shared server S000 is holding a lock on a resource.
C. The user assigned to shared servers S001 and S003 is waiting for a lock resource.
D. The user assigned to the shared server S002 is waiting for a lock resource to be received.

Answer: A

Question: 100
Which task can be accomplished by using the ALTER INDEX....COALESCE or REBUILD command?

A. Eliminating duplicate key values
B. Making an unusable index usable
C. Eliminating all blocks above high-water mark
D. Eliminating some or all the sparsely populated blocks

Answer: D

Question: 101
You executed the following query to view the status of a materialized view:

SQL> SELECT mview_name,rewrite_capability,refresh_method FROM user_mviews;
MVIEW_NAME REWRITE_CAPABILITY REFRESH_METHOD
----------------------------- ----------------------------- -------------------------
ITMV TEXTMATCH FORCE
Which statement is true regarding the rewrite and refresh mechanism?

A. The query of the materialized view contains restrictions on the use of query rewrite and the database performs a fast refresh if possible; otherwise, it does a complete refresh.
B. The query of the materialized view contains restrictions on the use of query rewrite and the database performs an incremental refresh.
C. The database can apply any rewrite rule that is supported and the materialized view is completely refreshed from the master table.
D. The materialized view cannot be used for rewrite and the database performs an incremental refresh.

Answer: A

Question: 102
You configured your database to run in Shared Server mode. You did not configure the large pool by using the LARGE_POOL_SIZE parameter. Which memory component would be used to store users' session information?

A. The Java Pool
B. The Shared Pool
C. The Streams Buffer
D. The Database Buffer Cache
E. The Program Global Area (PGA)

Answer: B

Question: 103
You moved your test database to the production environment. As a performance measure, you want to create a performance baseline for the production database that would have general performance statistics, SQL plans and usage, segment-level statistics, and parent and child latches. Which option would you use to create the baseline?

A. Use the TKPROF utility
B. Set the SQL_TRACE parameter to True
C. Use the Statspack with the highest snap level
D. Set the TIMED_STATISTICS parameter to True
E. Run the UTLBSTAT.SQL and UTLESTAT.SQL scripts

Answer: C

Question: 104
Which two statements are correct regarding stored outlines? (Choose two.)

A. The outlines are stored in the SYS schema.
B. The stored outlines always use the cost-based optimizer.
C. You cannot create several stored outlines for a single SQL statement.
D. The USE_STORED_OUTLINES parameter cannot be set at the session level.
E. When you set USE_STORED_OUTLINES to false and CREATE_STORED_OUTLINES to true, the Oracle database creates outlines but does not use them.

Answer: B, E

Question: 105
What provides the information about the free extents in a temporary tablespace before a sort operation?

A. LRU chain
B. sort extent pool
C. resource free list
D. free extent table (FET$)
E. used extent table (UET$)

Answer: B

Question: 106
You execute the following command as part of copying statistics to another database:

SQL> EXECUTE

DBMS_STATS.EXPORT_TABLE_STATS('SCOTT','EMP',NULL,'STATS',NULL,TRUE);

Which two options are correct recommendations that must be implemented before performing this step? (Choose two.)

A. The table that will hold the statistics must be analyzed.
B. The EMP table should be analyzed to obtain recent statistics.
C. The schema statistics for the user SCOTT should be collected.
D. The old statistics must be imported to the source database data dictionary.
E. The AVG_ROW_LEN and NUM_ROWS columns in the USER_TABLES view should be checked that the values are nonzero for EMP table.

Answer: B, E

Question: 107
View the Exhibit and examine the steps to create a database resource plan. Execution of the following procedure resulted in an error:

SQL> EXECUTE dbms_resource_manager.validate_pending_area();

What could be the reason for this?

```
SQL> EXECUTE dbms_resource_manager.create_pending_area();

PL/SQL procedure successfully completed.

SQL> EXECUTE dbms_resource_manager.create_consumer_group(consumer_group => 'OLTP',-
>  comment => 'Online users');

PL/SQL procedure successfully completed.

SQL> EXECUTE dbms_resource_manager.create_plan (plan => 'PRIUSERS',-
>  comment => 'DSS/Batch priority, ...' );

PL/SQL procedure successfully completed.

SQL> EXECUTE dbms_resource_manager.create_plan_directive (plan => 'PRIUSERS',-
>  group_or_subplan => 'OLTP',comment => 'Onine Group',CPU_P1 => 60);

PL/SQL procedure successfully completed.|
```

A. The ADMINISTER_RESOURCE_MANAGER privilege is not granted to the user.
B. The procedure should be executed before creating the resource plan directive.
C. The SYS_GROUP resource consumer group is not included in the resource plan directive.
D. The OTHER_GROUPS resource consumer group is not included in the resource plan directive.

Answer: D

Question: 108
In your database, you executed the following command to monitor the ind1 index:

SQL> SELECT lf_rows,del_lf_rows,del_lf_rows_len,lf_rows_len FROM INDEX_STATS
2 WHERE NAME = 'IND1';
LF_ROWS DEL_LF_ROWS DEL_LF_ROWS_LEN LF_ROWS_LEN
-------------- --------------------- --------------------------- --------------------
14 0 0 202

After this, the table is updated by a large transaction. Now, you want to check the index statistics. What would you do to update the INDEX_STATS view with the latest statistics?

A. Analyze the index with the ESTIMATE STATISTICS option
B. Analyze the index with the COMPUTE STATISTICS option
C. Analyze the index with the VALIDATE STRUCTURE option
D. Gather the table statistics first, and then analyze the index with any one of the options

Answer: C

Question: 109
Users in your production database complain that DMLs on an application table, ORDERS, is slow. You suspect the locking to be causing the contention and queried the V$LOCK dynamic performance view. While querying, you observe that the table is locked with lock type as TM but currently, there are no transactions or queries that are accessing the table. What could have caused the table to be locked?

A. The rows of a child table are locked.
B. The associated indexes are being analyzed.
C. The table is placed in the SYSTEM tablespace.
D. The structure of a child table is being modified.

Answer: B

Question: 110
Which statement is valid regarding the setting of the PCTFREE value for an index?

A. PCTFREE has no role to play when segment space management is automatic.
B. It keeps space reserved in a block for possible updates to existing rows in that block.
C. PCTFREE decides how much of a block to fill when the index is created or during a block split.
D. PCTFREE is the minimum percentage of a block that can be used for row data and overhead before new rows are added to the block.

Answer: C

Question: 111
You have online application users who are interested in seeing a few initial results of a query. Therefore you require the query to produce a few initial rows quickly regardless of the presence of statistics. Which optimizer mode would you choose?

A. RULE
B. CHOOSE
C. ALL_ROWS
D. FIRST_ROWS_n

Answer: D

Question: 112
A query on your table requires full table scan to fetch the records. What happens during this full table scan?

A. The Oracle server reads all the blocks below the high water mark.
B. The Oracle server reads all the blocks above the high water mark.
C. The Oracle server reads all the blocks in the segment irrespective of the high water mark.

D. The Oracle server reads only those blocks that contain data and does not read any empty blocks.

Answer: A

Question: 113
You are monitoring the redo log buffer. How would you check whether the user processes have to wait for the space in the redo log buffer?

A. by checking the V$WAITSTAT view for the "log buffer space" class
B. by checking the "redo buffer allocation retries" statistic in the V$SYSSTAT view
C. by checking the alert log file for the message "CHECKPOINT NOT COMPLETE"
D. by checking the "Log File Switch Completion" event in the V$SYSTEM_EVENT view

Answer: B

Question: 114
View the Exhibit and examine the TKPROF output. Which is the correct action required to improve the performance of the query?

```
select *
from
 cust   where id=50


call      count       cpu     elapsed       disk      query    current        rows
------- ------  -------- ----------- ---------- ---------- ---------- ----------
Parse        1      0.01        0.06          0          1          0           0
Execute      1      0.00        0.00          0          0          0           0
Fetch        2      0.24        0.55       1453       1797          0           3
------- ------  -------- ----------- ---------- ---------- ---------- ----------
total        4      0.25        0.62       1453       1798          0           3
```

A. Use the ALL_ROWS hint in the query
B. Increase the size of the database buffer cache
C. Decrease the value for the DB_FILE_MULTIBLOCK_READ_COUNT parameter
D. Create an index on the ID column or rebuild it if the existing index is not being used

Answer: D

Question: 115
While studying the statspack report, you observe a contention for the Sort Extent Pool latch, caused due to too many concurrent sorts. What action would you take to minimize the contention?

A. Modify the temporary tablespace to decrease the extent size
B. Resize the data file to increase the size of the temporary tablespace
C. Decrease the value of the SORT_AREA_RETAINED_SIZE parameter
D. Increase the PGA_AGGREGATE_TARGET to make more sorts stay in memory

Answer: D

Question: 116

ABC Corp. has an online transaction processing (OLTP) system that has just gone into production. The system was rigorously tested during the development phases. What is the first recommended activity that the DBA should perform before starting to monitor the performance of the system on a regular basis?

A. Running SQL*Trace
B. Collecting baseline statistics
C. Analyzing indexes and user tables
D. Arranging meetings with users once each two weeks to discuss performance

Answer: B

Question: 117
What would you do to ensure that important tables are not dropped in the middle of a long-running query?

A. Manually lock the tables in row share mode before the query starts executing
B. Decrease the value of the DML_LOCKS parameter before the query starts executing
C. Change the value of the ROW_LOCKING parameter to INTENT before the query starts executing
D. Decrease the value of the ENQUEUE_RESOURCES parameter before the query starts executing

Answer: A

Question: 118
You detected a sudden file I/O load on your database. Consequently, you need to isolate user sessions that are contributing to the load. What is the first diagnostic step that you can take in the given scenario?

A. query the V$SYSSTAT dynamic performance view to identify the SQL statements causing maximum I/O
B. use Top Session diagnostic pack in Oracle Enterprise Manager to identify the sessions that are contributing most to the I/O
C. use the findings of the Reorg Wizard to identify the SQL statements that are causing maximum I/O
D. query the V$SESSION dynamic performance view to check the value for the ROW_WAIT_ROW# column

Answer: B

Question: 119
You notice that large number of rows are migrated in one of the production tables. Which configuration could have caused this problem?

A. The default block size is very small.
B. PCTFREE is set to a very low value.
C. PCTUSED is set to a very high value.
D. The segment space management is set to AUTO.
E. The table is stored in a locally managed tablespace.

Answer: B

Question: 120

In an online transaction processing (OLTP) system, you notice that the transactions are waiting. Which two techniques would you use to determine if it is a locking problem? (Choose two.)

A. query V$SESSION to determine the sessions waiting for locks
B. query V$LOCKED_OBJECT to diagnose the deadlocked sessions
C. execute the utllockt.sql script to find the list of locking and waiting transactions
D. use the catblock.sql script to populate DBA_BLOCKERS and DBA_WAITERS views
E. query the DBA_WAITERS view to determine the sessions that are waiting for a locked resource

Answer: C, E

Question: 121
The following are the tasks that you need to perform to create a resource plan:

A) Creating resource consumer groups
B) Validating pending area
C) Creating a pending area
D) Creating resource plans
E) Creating plan directives
F) Submitting pending area

In which order would you perform these tasks?

A. C, D, A, E, B, F
B. C, B, F, E, A, D
C. C, B, E, D, A, F
D. C, F, E, A, D, B

Answer: A

Question: 122
The following statement that is used to create an index-organized table (IOT) failed with an error:

create table orders(
*
ERROR at line 1:
ORA-01450: maximum key length (3215) exceeded

What could be the reason?

A. A foreign key on another table is too long.
B. The primary key requires more than 3,215 bytes of space.
C. More than 3,215 columns are included in the composite Primary Key.
D. More than 3,215 bytes are needed for the aggregate of all the primary keys in the table.

Answer: B

Question: 123
You are working on the performance tuning of your database. What kind of performance tuning-related information would an alert log file provide? (Choose all that apply.)

A. Latch contention information
B. The time taken to archive logs
C. SQL statements that are consuming maximum resources

D. The top five sessions that are consuming maximum resources
E. Values of the non-default initialization parameters when the instance was started

Answer: B, E

Question: 124
The USER1 user executed the following command to create a private outline:

CREATE OR REPLACE PRIVATE OUTLINE outln1
ON SELECT COUNT(*) FROM SALES;
USER1 received the following error:
ORA-18009: one or more outline system tables do not exist

How should the error be resolved?

A. USER1 should create plan_table by using the utlxplan.sql script.
B. The CREATE_STORED_OUTLINE initialization parameter should be set to true at the session
   level.
C. The OUTLN user should execute the dbms_outln_edit.create_edit_tables procedure to create
   edit tables.
D. USER1 should execute the ANALYZE TABLE sales COMPUTE STATISTICS command
   before creating the stored outline.
E. USER1 should execute the dbms_outln_edit.create_edit_tables procedure to create edit tables
   in the USER1 schema.

Answer: E

Question: 125
You find that there is a contention in the free list of the ORDERS table. Which two tasks would
you accomplish to resolve the free list contention for the table? (Choose two.)

A. Reset the high water mark of the table
B. Increase the number of free lists of the table
C. Move the table to a locally managed tablespace
D. Export and import the table by using the Oracle Export and Import utility
E. Move the table to a tablespace that is using Automatic Segment Space Management

Answer: B, E

Question: 126
View the Exhibit and examine the Buffer Pool Advisory section in the statspack report. What size
would you recommend for the database buffer cache?

```
Buffer Pool Advisory for DB: ORCL   Instance: orcl   End Snap: 8
-> Only rows with estimated physical reads >0 are displayed
-> ordered by Block Size, Buffers For Estimate

          Size for  Size      Buffers for  Est Physical          Estimated
P    Estimate (M) Factr        Estimate   Read Factor      Physical Reads
---  ------------ -----    -------------  ------------    ----------------
D              8   .5           1,001          3.49               487,871
D             16  1.0           2,002          1.00               139,786
D             24  1.5           3,003          0.48               100,258
D             32  2.0           4,004          0.48                66,505
D             40  2.5           5,005          0.48                66,505
D             48  3.0           6,006          0.48                66,505
D             56  3.5           7,007          0.48                66,505
D             64  4.0           8,008          0.48                66,505
D             72  4.5           9,009          0.48                66,505
D             80  5.0          10,010          0.48                66,505
D             88  5.5          11,011          0.48                66,505
D             96  6.0          12,012          0.48                66,505
D            104  6.5          13,013          0.48                66,505
D            112  7.0          14,014          0.48                66,505
D            120  7.5          15,015          0.48                66,505
D            128  8.0          16,016          0.48                66,505
D            136  8.5          17,017          0.48                66,505
D            144  9.0          18,018          0.48                66,505
D            152  9.5          19,019          0.48                66,505
D            160 10.0          20,020          0.48                66,505
     ---------------------------------------------------------------------
```

A. 32 MB
B. 24 MB
C. 16 MB
D. 40 MB

Answer: A

Question: 127
You plan to have a small number of large extents instead of large number of small extents for the objects in a dictionary managed tablespace. What could be the reason for this decision?

A. It minimizes space wastage in the database.
B. It can reduce the frequency of dynamic extent allocation.
C. It is easier for the Oracle server to arrange blocks for the extents.
D. It improves performance in the case when indexes are used for searching.

Answer: B

Question: 128
Which two statements are true regarding the User Global Area (UGA) for a database that is running in Shared Server mode? (Choose two.)

A. The Stack space is a part of the UGA.
B. The UGA is always stored in the Shared Pool.
C. The UGA is stored in the Shared Pool if the large pool is not configured.
D. All the user information is stored in the UGA, therefore individual shared servers do not need
    to remain bound to a user session.

E. Any of the shared server processes can access users' session information, which is stored in the UGA.

Answer: C, E

Question: 129
View the Exhibit.
You issued the query on the V$BUFFER_POOL_STATISTICS dynamic performance view twice.
What could the reasons for no increase in the BUF_GOT value? (Choose all that apply.)

```
SQL> SELECT name, buf_got FROM v$buffer_pool_statistics WHERE name='KEEP';

NAME                      BUF_GOT
-------------------- -----------
KEEP                        819284

SQL> SELECT count (*) FROM hr.departments;

  COUNT(*)
----------
        27

SQL>  SELECT name, buf_got FROM v$buffer_pool_statistics WHERE name='KEEP';


NAME                      BUF_GOT
-------------------- -----------
KEEP                        819284
```

A. The keep pool is full.
B. The count(*) function bypasses the keep pool.
C. The query is being implicitly rewritten using a materialized view.
D. The object you are querying has its blocks already present in the buffer cache.
E. The object you are querying has blocks already present in default or recycle pool.

Answer: D, E

Question: 130
What is the effect of setting the FAST_START_MTTR_TARGET initialization parameter?

A. A higher value for the FAST_START_MTTR_TARGET initialization parameter results in less time spent on reading redo during instance startup.
B. A lower value for the FAST_START_MTTR_TARGET initialization parameter results in more time spent on reading redo during instance startup.
C. A lower value for the FAST_START_MTTR_TARGET initialization parameter results in an increase in the frequency of the DBWRn process writing dirty buffers.
D. A higher value for the FAST_START_MTTR_TARGET initialization parameter results in the LGWR process writing redo log buffers faster.

Answer: C

Question: 131
View the Exhibit and examine the commands.
What could be the reason for the INDEX_STATS view not being updated?

A. The index is not of the B-tree type.
B. The table has to be analyzed first.
C. The table statistics have not been gathered.
D. The index has to be analyzed with the COMPUTE STATISTICS option.
E. The index has to be analyzed with the VALIDATE STRUCTURE option.

Answer: E

Question: 132
Which scenario would cause contention in the free list of a segment?

A. Large number of inserts being performed by a single user session
B. Large number of deletes being performed by a single user session
C. Large number of inserts being performed by multiple user sessions
D. Large number of updates being performed by multiple user sessions

Answer: C

Question: 133
View the Exhibit.
Check the following configuration:

NAME VALUE
-------------------------------------- -------------
query_rewrite_enabled TRUE
query_rewrite_integrity enforced
optimizer_mode CHOOSE
You executed the following command on SALES and TIMES tables:
SELECT s.prod_id
, t.fiscal_month_number
, sum(s.amount_sold) AS sum_amount
FROM sales s, times t
WHERE s.time_id = t.time_id
AND t.fiscal_year = 2000
GROUP BY s.prod_id, t.fiscal_month_number
ORDER BY s.prod_id, t.fiscal_month_number;

You observed from the execution plan that the query is not rewritten on SALES_PROD. What could be the reason for this?

```
SQL> create materialized view sales_prod
  2          build immediate
  3          enable query rewrite
  4  as
  5  SELECT s.prod_id
  6  ,       t.fiscal_month_number
  7  ,       sum(s.amount_sold) AS sum_amount
  8  FROM    sales s, times t
  9  WHERE   s.time_id = t.time_id
 10  AND     t.fiscal_year = 2000
 11  GROUP   BY s.prod_id,
t.fiscal_month_number;

Materialized view created.
```

A. The OPTIMIZER_MODE is not set to RULE.
B. The QUERY_REWRITE_INTEGRITY is not set to TRUSTED.
C. The definition of materialized view does not have the ORDER BY clause.
D. The SALES table has been updated to make changes in the amount_sold column.

Answer: D

Question: 134
You logged in as sysdba and executed the catblock.sql script in your database. What does the script do?

A. Creates the dba_locks and dba_blockers views
B. Kills the sessions that are holding unnecessary locks
C. Displays the details about sessions that are waiting for locks
D. Displays the details about sessions that are blocking objects

Answer: A

Question: 135
What are the two reasons for creating a reverse key index on a column? (Choose two.)

A. When the column contains Boolean values
B. When the column implements an inverted list attribute
C. When the column is mainly used for value range scans
D. When the column is populated using sequential numbers
E. When you want to allow insertions to be distributed across all the leaf blocks in the index

Answer: D, E

Question: 136
You notice that queries such as SELECT count(*) FROM ..on a couple of tables are consuming a lot of time. On further investigation, you find that a huge number of records from these tables have been removed but the high water mark has not been reset. Which operations on the tables would reset the high water mark and consequently improve performance? (Choose all that apply.)

A. Moving the tables to another tablespace
B. Moving the tables in the same tablespace
C. Altering the table to deallocate unused space
D. Exporting, truncating, and importing the tables

Answer: A, B, D

Question: 137
View the Exhibit and examine the parameters. Why would the DBA specify different values for the SESSIONS and SHARED_SERVER_SESSIONS parameters?

```
SQL> SHOW PARAMETER SESSION

NAME                               TYPE        VALUE
---------------------------------- ----------- --------------------------------
java_max_sessionspace_size         integer     0
java_soft_sessionspace_limit       integer     0
license_max_sessions               integer     0
license_sessions_warning           integer     0
logmnr_max_persistent_sessions     integer     1
mts_sessions                       integer     150
session_cached_cursors             integer     0
session_max_open_files             integer     10
sessions                           integer     170
shared_server_sessions             integer     150
```

A. To reserve sessions for the dispatcher processes
B. To reserve user sessions for dedicated connections
C. To reserve sessions for Oracle server's internal sessions
D. To reserve sessions for use when the database instance is hung

Answer: B

Question: 138
You want to reserve memory in your database instance to accommodate large memory allocations during operations such as PL/SQL and trigger compilation. What must you do to achieve this?

A. Configure Large Pool
B. Increase the size of the Shared Pool
C. Configure the User Global Area (UGA)
D. Configure the Shared Pool Reserved area
E. Increase the size of the System Global Area (SGA) by using the SGA_MAX_SIZE parameter

Answer: D

Question: 139
After analyzing an index segment, you find that the B-level is high and most of the index blocks contain very few entries due to a large number of DELETE operations. Which statement is true?

A. Index segment blocks with 25% deleted entries are put into the free list.
B. Index segment blocks with 75% deleted entries are put into the free list.
C. Index segment blocks with 50% deleted entries are put into the free list.
D. Index segment blocks with 100% deleted entries are put into the free list.

Answer: D

Question: 140
Which two activities would adversely affect database performance? (Choose two.)

A. Doing a soft parse
B. Performing archiving
C. Less frequent checkpointing
D. Defining multiple buffer pools
E. Taking online back up of data files

Answer: B, E

Question: 141
Identify the common causes of poor library cache performance. (Choose all that apply.)

A. Large pool too small
B. Shared pool too small
C. SQL cursor not being shared
D. User Global Area (UGA) too small
E. High number of invalidations occurring

Answer: B, C, E

Question: 142
What is the purpose of executing the UTLCHAIN.SQL script in your database?

A. To analyze a table for listing the chained rows
B. To create a table for holding the information about chained rows
C. To collect and store the statistics of a table containing chained rows
D. To collect and store the statistics of a table containing migrated rows

Answer: B

Question: 143
You have objects that have a variable reference pattern. These are referenced vigorously for some time, and then not referenced at all for a long time. As a result, these objects are being aged out from Oracle memory buffers frequently. Which two strategies would you recommend to avoid aging out of erratically accessed objects? (Choose two.)

A. Increasing PGA size
B. Increasing large pool size
C. Increasing buffer cache size
D. Creating RECYCLE cache using the DB_RECYCLE_CACHE_SIZE parameter
E. Creating a KEEP cache using the DB_KEEP_CACHE_SIZE parameter and altering the objects to that cache

Answer: C, E

Question: 144
Your production environment has the following features:

- The database is running in ARCHIVELOG mode.
- There are two online redo log groups.

Exam Name:	**Oracle 9i Database: Performance Tuning**		
Exam Type:	**Oracle**		
Exam Code:	**1Z0-033**	Total Questions:	**165**

- The redo log files are not multiplexed.
- The redo log files are frequently overwritten because of a large number of transactions.

What is the effect on the database if one of the online redo log files is not yet archived, and the log writer (LGWR) process attempts to overwrite it because of a log switch?

A. The transactions halt until the redo log file becomes available or is archived.
B. The online redo log file is not archived and database operations continue as normal.
C. The database instance is shut down immediately, requiring an instance recovery to be performed.
D. The number of the Archiver (ARCn) processes is dynamically increased by the RDBMS for faster archiving of the online redo log file.
E. The database instance is shut down immediately, requiring a media recovery to be performed.

Answer: A

Question: 145
You are working with an application that has the following requirements:

- Fast key-based access to table data
- Lots of queries involving exact match and range search
- Retrieval of data based on primary key

Which is the correct storage structure for this purpose?

A. Cluster
B. Heap table
C. Partitioned table
D. Index-organized table

Answer: D

Question: 146
You executed the following commands to view the statistics:

SQL> SELECT table_name,blocks,num_rows FROM dba_tables WHERE table_name='CUST'
and
owner='SCOTT';
TABLE_NAME BLOCKS NUM_ROWS
-------------------- ------------ ---------------
CUST 399 8192
SQL> SELECT index_name,blevel,leaf_blocks,clustering_factor FROM dba_indexes WHERE
index_name='CUIND' and owner='SCOTT';
INDEX_NAME BLEVEL LEAF_BLOCKS CLUSTERING_FACTOR
-------------------- ---------- -------------------- ------------------------------
CUIND 2 50 7946

Which statement is true regarding the output?

A. The index is a well-ordered index.
B. Data blocks are less likely to be revisited.
C. The index entries randomly point at different data blocks.
D. The number of data blocks that will be accessed while scanning the index is 399.

Answer: C

Question: 147

View the Exhibit and examine the output to check the resource plan directives associated with the groups in a plan. Which three statements are true regarding the output? (Choose three.)

```
select plan, group_or_subplan, type, cpu_p1, cpu_p2, cpu_p3, cpu_p4, status
 from dba_rsrc_plan_directives
order by 1,2,3,4,5,6;

PLAN                   GROUP_OR_SUBPLAN     TYPE            CPU_P1     CPU_P2     CPU_P3     CPU_P4     STATUS
---------------------  -------------------  --------------  ---------- ---------- ---------- ---------- -------
SINGLE_LEVEL_PLAN      ADHOC_GROUP          CONSUMER_GROUP        10          0          0          0 ACTIVE
SINGLE_LEVEL_PLAN      BATCH_GROUP          CONSUMER_GROUP        10          0          0          0 ACTIVE
SINGLE_LEVEL_PLAN      OLTP_GROUP           CONSUMER_GROUP        80          0          0          0 ACTIVE
SINGLE_LEVEL_PLAN      OTHER_GROUPS         CONSUMER_GROUP         0          0          0          0 ACTIVE
4 rows selected.
```

A. The users belonging to OTHER_GROUPS will never be able to run a query.
B. Even if there is only one active session running, that session is allowed to utilize all the available CPU resources according to the directives of the plan.
C. It is possible to raise the priority for the online transaction processing (OLTP) group from 80% to 90% without changing the priority for any other group.
D. In a fully loaded system, sessions for OTHER_GROUPS have to wait for other sessions to finish work if all CPU is used by first three groups.
E. In a heavily used system, the sessions in the OLTP group have the ability to consume 80% or more of the total CPU resources available.

Answer: B, D, E

Question: 148

You executed a series of UPDATE statements on one of the tables. This resulted in row migration in the table. You want to eliminate row migration. Which action on the table would eliminate row migration? (Choose all that apply.)

A. Export, drop and importing the table
B. Executing the ALTER TABLE..MOVE command on the table
C. Allocating an extent to the table by using the ALTER TABLE command
D. Executing the ALTER TABLE..DEALLOCATE UNUSED command on the table

Answer: A, B

Question: 149

You want SMITH to create database resource plans using database resource manager. To enable this, you decide to grant the ADMINISTER_RESOURCE_MANAGER privilege to SMITH. Which statement regarding this privilege is true?

A. This privilege cannot be granted to a role.
B. This privilege is granted through the SQL GRANT command.
C. This privilege is granted through the dbms_resource_manager package.
D. This privilege cannot be granted to SMITH; only SYS can have this privilege.
E. This privilege is granted through the dbms_resource_manager_privs package.

Answer: E

Question: 150

You executed the following command to create an index on IPCOST column of PRODTB table:

SQL> CREATE INDEX ipcost ON prodtb(pcost) NOSORT;

Why would it be useful to create the index with the NOSORT clause?

A. It uses one of the existing indexes for sorting.
B. It creates the index with sorting on the basis of the primary key.
C. It creates an invalid index that needs to be rebuilt after values in the column are sorted.
D. It cuts down the time needed to build the index and workareas space and temp tablespace I/O needed to build the index.

Answer: D

Question: 151
Which two statements describe the correct guidelines for setting the extent size for a temporary tablespace to achieve better I/O performance? (Choose two.)

A. Never use the UNIFORM option in case of a locally managed temporary tablespace.
B. The UNIFORM option should not be used with a locally managed temporary tablespace when PGA_AGGREGATE_TARGET is set to a nonzero value.
C. The INITIAL and NEXT values should be an integer multiple of SORT_AREA_SIZE in case of a dictionary-managed temporary tablespace.
D. The PCTINCREASE should be set to zero in case of a dictionary-managed temporary tablespace.

Answer: C, D

Question: 152
You executed the following statement in session 1:

SQL> SELECT ename FROM emp WHERE deptno=10 FOR UPDATE;
ENAME

CLARK
KING
MILLER
In session 2, you executed the following statement on the same table:
SQL> UPDATE emp SET sal=sal*1.2 WHERE deptno=30;

Which statement regarding UPDATE in session 2 is correct if the database is using the default locking mechanism?

A. This updates the table because the query above does not hold any locks.
B. This does not update the table because the table-level lock is held by the first session.
C. This updates the table because you are updating a different column than that selected in the first session.
D. It updates the rows in department 30 because the three locked rows are rows in department 10. As the default is row level locking the rows are updated.

Answer: D

Question: 153
Which view would you query to monitor cumulative total waits for all events and all sessions?

A. V$SYSSTAT
B. V$SYSTEM_EVENT
C. V$SESSION_WAIT

D. V$SYSTEM_STATUS

Answer: B

Question: 154
You are running low on a system resource, such as read/write throughput or CPU cycle. Which technique would you adopt to detect the most resource-consuming statements in your application code for a specific session?

A. Use the contents of the alert log file
B. Use the report generated by UTLBSTAT/UTLESTAT
C. Take a trace and use TKPROF to see the formatted output
D. Use the STATSPACK output generated with snapshot level 0

Answer: C

Question: 155
View the Exhibit and examine the statistics of the CUST table. Which statement is true regarding optimization of space usage by the CUST table?

```
SQL> select count(distinct dbms_rowid.rowid_block_number(rowid)) from cust;

COUNT(DISTINCTDBMS_ROWID.ROWID_BLOCK_NUMBER(ROWID))
---------------------------------------------------
                                                 23

SQL> select segment_name,blocks from user_segments where segment_name='CUST';

SEGMENT_NAME         BLOCKS
--------------- ----------
CUST                    528

SQL> select table_name,blocks,empty_blocks from user_tables
  2  where table_name='CUST';

TABLE_NAME                         BLOCKS EMPTY_BLOCKS
------------------------------- ---------- ------------
CUST                                  399          128
```

A. Use the ALTER TABLE....DEALLOCATE UNUSED command because a large number of
   blocks are not used below high-water mark.
B. Use the ALTER TABLE....MOVE command because a large number of blocks are not used
   below high-water mark.
C. Use the DBMS_SPACE.FREE_BLOCKS procedure because a large number of blocks are not
   used below high-water mark.
D. Use the DBMS_SPACE.UNUSED_SPACE procedure because a large number of blocks are
   not used below high-water mark.

Answer: B

Question: 156
Identify three conditions that cause the DBWn process to write blocks from the database buffer cache to data files. (Choose three.)

A. When a segment is dropped

B. When the redo log buffer is one-third full
C. When a tablespace is taken offline normally
D. When the buffer search threshold in the database buffer cache is exceeded
E. When the requested blocks are available in the database buffer cache but not in data files

Answer: A, C, D

Question: 157
You are working on a test database. The DML monitoring feature is enabled for all the tables that support monitoring. You executed the following command to gather statistics:

exec dbms_stats.gather_schema_stats(ownname => NULL, options => 'GATHER AUTO');

Which statement is correct regarding the outcome of this command?

A. It gathers statistics for all user objects in the database.
B. It implicitly determines which objects need new statistics and gathers all necessary statistics automatically.
C. It returns a list of objects with stale statistics and automatically gathers statistics on those objects only.
D. It returns a list of objects with no statistics and automatically gathers statistics on those objects only.

Answer: B

Question: 158

You receive complaints from database users about performance degradation. On investigation, you find that this is due to unnecessary sort operations in the database. Which are the two operations that would reduce sorting? (Choose two.)

A. Using sort-merge join instead of the nested-loop join
B. Using UNION instead of UNION ALL, wherever possible
C. Creating indexes on the columns that are used in equijoins
D. Using ANALYZE command for large objects with ESTIMATE clause instead of COMPUTE

Answer: C, D

Question: 159
The HR team is complaining about the poor response time of SQL statements on certain tables. On investigation, you find that the generation of recursive calls (in addition to SQL statements) to perform space management operations is the cause for the poor response time. To solve the problem, you executed the following command as the first step:

SQL> SELECT owner, table_name, blocks, empty_blocks
2 FROM dba_tables
3 WHERE empty_blocks/(blocks+empty_blocks)<.1;
OWNER TABLE_NAME BLOCKS EMPTY_BLOCKS
--------------- ----------------- ------------ ---------------------
HR REGIONS 1205 30
HR LOCATIONS 110 5

Select the next step that you would perform.

A. Allocating extent to the tables by using the ALTER TABLE command

B. Changing the block size by using the ALTER TABLESPACE command
C. Exporting, dropping, and importing the REGIONS and LOCATIONS tables
D. Creating the necessary indexes on the REGIONS and LOCATIONS tables

Answer: A

Question: 160
When tuning a database server, which three goals should you focus on? (Choose three.)

A. Minimize response time for OLTP users
B. Maximize throughput for batch processing
C. Minimize response time for batch processing
D. Increase load capability for the decision-support systems
E. Minimize parallelism for online transaction processing (OLTP) systems

ANSWER: A, B, D

Question: 161
Which three statements describe good practices for an application developer to reduce locking problems? (Choose three.)

A. Avoid coding long-running transactions.
B. Let the database handle locks in default locking mode.
C. Always explicitly code the locks as per the requirement of the application.
D. Allow escalation of row locks to block locks if too many row locks cause a problem.
E. Long-running transactions that interfere with other sessions should be run during off-peak
   hours.

Answer: A, B, E

Question: 162
You queried the V$SYSTEM_EVENT view and noticed that there are waits for the "Log File Switch Completion" event. What could be the reason for these waits?

A. The number of ARCn processes is not sufficient.
B. The number of DBWRn processes is not sufficient.
C. There is disk I/O contention on the online redo log files.
D. The value set for the FAST_MTTR_TARGET initialization parameter is too high.

Answer: C

Question: 163
In an effort to tune the OLTP database, you decide to increase the value set for the PGA_AGGREGATE_TARGET parameter. What could be the reason for this action?

A. In your database, the ratio of disk sorts to memory sorts is less than 5%.
B. In your database, the ratio of disk sorts to memory sorts is more than 20%.
C. In your database, the sort segment in the temporary tablespace is very small.
D. In your database, you are getting many hash joins instead of sort merge joins.

Answer: B

Question: 164
Which three types of objects can you keep in the shared pool by using the dbms_shared_pool.keep procedure? (Choose three.)

A. Views
B. Tables
C. Triggers
D. Sequences
E. PL/SQL functions

Answer: C, D, E

Question: 165
Jones and a couple of other users complain that their transactions on one of the application tables, TECH, are waiting for a response. On investigation, you find that one of the users, Smith, has not committed his transaction on the TECH table and he is not at his desk. What would you do to release the lock irrespective of the end result of Smith's transaction?

A. Kill Smith's session
B. Perform a rollback operation in your session on behalf of Smith's session
C. Smith's session will be terminated automatically by altering the user SMITH to use a profile with an idle limit.
D. No explicit action required (The lock will be released by the Oracle database automatically after specified number of seconds in the DML_LOCKS parameter.)

Answer: A

End of Document