



Exam : 1Z0-033

Title : Performance Tuning

Ver : 08.25.06

QUESTION 1:

Which two statements regarding OLTP systems are true? (Choose two)

- A. Use literals for optimally shared SQL rather than bind variables to keep the overhead of parsing to a minimum.
- B. To avoid the performance load of dynamic space allocation, allocate space explicitly so tables, clusters and indexes.
- C. B-tree indexing is preferred to bitmap indexing, because of locking issues affecting DML operations.
- D. Use hash clusters especially on tables that are heavily inserted into, because of the use of space and the number of blocks that need to be visited.
- E. Use application code to enforce rules instead of constraints, because constraints are extremely expensive to process.

Answer: B, C

Explanation:

Online Transaction Processing (OLTP) systems tend to be accessed by large numbers of users doing short DML transactions. Users of OLTP systems are primarily concerned with throughput: the total time it takes to place an order, remove an item from inventory, or schedule an appointment. To avoid the performance load of dynamic space allocation, you need to allocate space explicitly so tables, clusters and indexes. Bitmap indexes will not work good for the OLTP systems because of locking issues affecting DML operations. B-tree indexes can handle this easier and effectively.

Incorrect Answers

A: You cannot use literals for optimally shared SQL rather than bind variables because of nature of OLTP systems: they work effectively using bind variables because of performing DML operations.

D: Hash clusters work more effective for DSS (Decision Support Systems): high level inserts and updates will eliminate the advantage of hash clusters which require lower level of DML activity as DSS systems provide.

E: Constraints are not expensive to process: they need to be used to avoid additional application code creation to enforce a business rules.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 187-188
Chapter 3: SQL Application Tuning and Design

QUESTION 2:

When performing a sort operation, you notice that there are a large number of sorts requiring I/O to the disk. Which parameter could be increased to allow more sorts to be performed in memory?

- A. SORT_AREA_SIZE
- B. LARGE_POOL_SIZE

- C. SORT_AREA_RETAINED_SIZE
- D. SORT_MULTIBLOCK_READ_COUNT

Answer: A

Explanation:

The amount of memory set aside for each user's Server Process to perform these sort operations is impacted by the following init.ora parameters: SORT_AREA_SIZE, SORT_AREA_RETAINED_SIZE, PGA_AGGREGATE_TARGET, WORKAREA_SIZE_POLICY. The default value for SORT_AREA_SIZE is OS-dependent. The minimum size for this parameter is equivalent to six Oracle blocks. The maximum size is OS-dependent. SORT_AREA_SIZE specifies how much memory each user's Server Process should set aside to perform in-memory sort operations.

Incorrect Answers

B: LARGE_POOL_SIZE parameter is used to buffer I/O server processes as well as backup and recovery. It also caches session data when the Shared Server feature is used.

C: Once a sort operation is complete, if the sort area still contains sorted rows that need to be returned to the user, the user's Server Process reduces the memory set aside for the final fetch to the value specified by SORT_AREA_RETAINED_SIZE.

D: SORT_MULTIBLOCK_READ_COUNT specifies the number of database blocks to read each time a sort performs a read from a temporary segment. Temporary segments are used by a sort when the data does not fit in SORT_AREA_SIZE of memory. In these situations, sort writes out sections of data to temporary segments in the form of sorted runs. Once all the data has been partially sorted to these runs, sort merges the runs by reading pieces of them from the temporary segment into memory to produce the final sorted output. If SORT_AREA_SIZE is not large enough to merge all the runs at once, subsets of the runs are merged in a number of merge passes.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 412-413
Chapter 8: Tuning Disk I/O

QUESTION 3:

Which statement could require a sort?

- A. SELECT DISTINCT dept_id
FROM emp;
- B. UPDATE emp SET salary=salary*1.1
WHERE id=7722;
- C. SELECT emp_id, name
FROM emp
WHERE emp-id=7722;
- D. SELECT emp_id, name
FROM emp
WHERE emp_id BETWEEN 7722 and 7100;

Answer: A

Explanation:

The types of SQL statements that can cause database sorts to occur include the following: ORDER BY, GROUP BY, SELECT DISTINCT, UNION, INTERSECT, MINUS, ANALYZE, CREATE INDEX, joins between tables on columns that are not indexed.

Incorrect Answers

B: This UPDATE statement does not require any sort operation.

C: This SELECT statement does not require any sort operation.

D: Usage of BETWEEN clause in the SELECT statement will not cause any sort operation.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 411-412

Chapter 8: Tuning Disk I/O

QUESTION 4:

Which two views can be used to detect lock contention? (Choose two)

A. V\$LOCK

B. V\$LOCKED_OBJECT

C. V\$LOCK_CONTENTION

Answer: A, B

Explanation:

When left to its default mechanisms, Oracle generally does a very effective job of managing locking. When it does occur, lock contention can be identified using the V\$LOCK and V\$LOCKED_OBJECT dynamic performance views, the DBA_WAITERS and DBA_BLOCKS data dictionary views, and the OEM Performance Manager GUI.

Incorrect Answers

C: There is no V\$LOCK_CONTENTION data dictionary view in Oracle.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 487-489

Chapter 9: Tuning Contention

QUESTION 5:

The database includes tables with static data, which are used for queries only. To which size should you set PCTFREE for this type of table?

A. 0

B. 50

C. 20

D. 10

Answer: A

Explanation:

1Z0-033

The parameter PCTFREE tells Oracle how much space to set aside in each block to store row updates. Since we have a database with static data, PCTFREE can be set to 0 to use the disk space more effectively.

Incorrect Answers

B: By setting this parameter to 50 you will preserve 50% of data block for the future inserts. It will be just space wasting because of static nature of data inside the tables.

C: This size still cannot be considered as appropriate for the database with static data.

D: You can use this size for the database with minimal level of update activity.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 401

Chapter 8: Tuning Disk I/O

QUESTION 6:

Which action could potentially cause checkpoints to take longer?

- A. Increasing the number of redo log groups.
- B. Increasing the size of rollback segments.
- C. Decreasing the value of the REDO_LOG_BUFFERS parameter.
- D. Increasing the value of the FAST_START_IO_TARGET parameter.

Answer: D

Explanation:

By increasing the FAST_START_IO_TARGET parameter you tell Oracle that it can use more time to perform recovery after instance crash. So it will cause that checkpoints will take longer. This parameter supersedes the FAST_START_MTTR_TARGET parameter if configured. The FAST_START_MTTR_TARGET init.ora parameter is used to specify a mean time (in seconds) to recover the instance following an instance failure.

Incorrect Answers

A: By increasing the number of redo log groups you will not decrease a time for checkpoints.

B: The size of rollback segments has nothing to do with checkpoints duration.

C: There is no REDO_LOG_BUFFERS parameter in Oracle.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 343-345

Chapter 7: Tuning Redo Mechanisms

QUESTION 7:

When a deadlock shutdown is detected by Oracle, where is the trace file generated?

- A. SQL_TRACE
- B. TRACE_DEST
- C. USER_DUMP_DEST
- D. CORE_DUMP_DEST
- E. BACKGROUND_DUMP_DEST

Answer: C

Explanation:

Oracle creates a user trace file when processing errors in a user's Server Process. A deadlock error will be shown in a user trace file which is located in the USER_DUMP_DEST. While events such as deadlocks automatically generate trace files, full scale tracing of user sessions does not occur unless the user or DBA requests it.

Incorrect Answers

A: SQL_TRACE is the main method for collecting SQL Execution information in Oracle collecting a wide range of information and statistics that can be used to tune SQL operations.

B: The TRACE_DEST parameter does not specify the location of this file.

D: CORE_DUMP_DEST should point to the directory where core dumps from the Oracle server will be placed. A core dump is a memory image of the Oracle shadow process produced when an unexpected, unrecoverable or invalid condition occurs. Note that Oracle should always try to write a trace file before producing a core dump.

E: BACKGROUND_DUMP_DEST specifies the pathname (directory or disc) where debugging trace files for the background processes (LGWR, DBWn, and so on) are written during Oracle operations.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 40-41
Chapter 2: Sources of Tuning Information

QUESTION 8:

Which two statements are true of the buffer cache? (Choose two)

- A. The buffer pools in Oracle9i can be resized dynamically.
- B. The blocks in the buffer cache are managed using one list.
- C. The blocks in the buffer cache are managed using two lists.
- D. In Oracle9i buffer caches can be dynamically added.

Answer: A, C

Explanation:

Beginning with Oracle9i, the dynamic SGA infrastructure will allow for the sizing of the Buffer Cache, Shared Pool and the Large Pool without having to shutdown the database. The buffer pools in Oracle9i can be resized dynamically without stopping instance with commands like ALTER SYSTEM SET DB_CACHE_SIZE=<size>. The blocks in the buffer cache are managed using two lists.

Incorrect Answers

B: The blocks in the buffer cache are managed using two lists, not one list.

D: Buffer caches, other than DEFAULT, i.e. DB_KEEP_CACHE_SIZE or DB_RECYCLE_CACHE_SIZE cannot be dynamically added. They need to be present in the init.ora or spfile.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 275-277
Chapter 5: Tuning the Database Buffer Cache

QUESTION 9:

For which reason would you query V\$SYSSTAT?

- A. Name of the sort segment.
- B. Free space available for a sort segment.
- C. Number of disk sorts performed since startup.
- D. Number of users active on individual sort segments.

Answer: C

Explanation:

Sort activity can be monitored using the V\$SYSSTAT and V\$SORT_SEGMENT dynamic data dictionary views, using the output from STATPACK and REPORT.TXT, and using the output from the OEM Performance Manager.

Incorrect Answers

- A: This view does not show a name of the sort segment.
- B: It is not used to show a free space available for a sort segment.
- D: V\$SYSSTAT does not provide information about number of users active on individual sort segments.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 414-415
Chapter 8: Tuning Disk I/O

QUESTION 10:

Which two statements about plan stability and stored outlines are true? (Choose two)

- A. You can group outlines in categories.
- B. You can only have one stored outline per SQL statement.
- C. Plan stability only wants when SQL statements match textually.
- D. Stored outlines are saved in the data dictionary (SYS schema).
- E. Stored outlines become invalid when you analyze the associated objects.

Answer: A, C

Explanation:

Oracle9i maintains predefined execution plans in the data dictionary in the form of stored outlines. Ascertain whether the statement you have identified can be lumped into two or more broad categories of statements. Through a process known as plan equivalence, a SQL statement must exactly match the original statement used to generate an outline in order for the stored outline to be used. Even an identical statement with the addition of a hint or comment will not be considered equivalent for stored outline usage.

Incorrect Answers

- B: You can have as many stored outline per SQL statement as you need.
- D: Stored outlines are saved in the OUTLN schema.

E: Stored outlines do not become invalid when you analyze the associated objects. That's why they are used to preserve predefined execution plans.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 131-140
Chapter 3: SQL Application Tuning and Design

QUESTION 11:

What does this statement do?

```
SQL.ANALYZE INDEX index_name VALIDITY STRUCTURE;
```

- A. It places information into the INDEX_STATS view and allows for the monitoring of space used by an index.
- B. It provides information in the INDEX_HISTOGRAM view to indicate whether an index is invalid or valid.
- C. It provides information in the DBA_INDEXES view for the COST BASED Optimizer when choosing an execution plan.
- D. It checks to see if each index entry points to a valid row in the underlying table and saves the results in the DBA_INDEXES view.
- E. It check to see of each row in the underlying table has an index entry and saves the results into the DBA_INDEXES view.

Answer: A

Explanation:

The VALIDATE STRUCTURE option for the ANALYZE command populates the data dictionary view INDEX_STATS with values.

Incorrect Answers

B: This command does not provide information in the INDEX_HISTOGRAM view.

C: This command does not provide information in the DBA_INDEX data dictionary view.

D:

E:

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 150

Chapter 3: SQL Application Tuning and Design

QUESTION 12:

Which three types of statistics are reported in report.txt after running UTLESTAT SQL?
(Choose three)

- A. Locking statistics.
- B. Memory usage statistics.
- C. Explain plan statistics.
- D. Library cache statistics.
- E. Buffer busy wait statistics.
- F. Rollback contention statistics.

Answer: D, E, F

Explanation:

There are three types of statistics reported in the REPORT.TXT file after running UTLESTAT.SQL: library cache statistics, buffer busy wait statistics, rollback contention statistics.

Incorrect Answers

A: Locking statistics are not reported in REPORT.TXT file.

B: There is no memory usage statistics in REPORT.TXT after running UTLESTAT.SQL.

C: Explain plan statistics are not reported in REPORT.TXT file.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 50-53
Chapter 2: Sources of Tuning Information

QUESTION 13:

What are two main OLTP requirements? (Choose two)

- A. Use bind variables rather than literals in your SQL code.
- B. Analyze your tables regularly to refresh optimizer statistics.
- C. Create multiple small rollback segments as opposed to a few big ones.
- D. Create indexes on all columns that are regularly used in query predicates.
- E. Set up appropriate default storage parameter values for dynamic (implicit) space allocation.

Answer: C, E

Explanation:

Due to high level of DML activity on the OLTP system you need to use more small rollback segments rather than few large rollback segments. To avoid the performance load of dynamic space allocation, you need to allocate space explicitly so tables, clusters and indexes.

Incorrect Answers

A: This one is not main requirement of the OLTP system.

B: Table and index statistics should be gathered regularly if the CBO is used because of data volumes tend to change quickly in OLTP systems. But this requirement cannot be considered as main OLTP requirement.

D: It is not recommended to create indexes on ALL columns that are regularly used in query predicates because of multiple INSERT and UPDATE operations will cause overload on the database trying to rebuild all indexes after each transaction.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 187-188
Chapter 3: SQL Application Tuning and Design

QUESTION 14:

Which tablespace is used as the temporary tablespace if 'TEMPORARY TABLESPACE' is not specified for a user?

- A. TEMP
- B. DATA
- C. SYSTEM
- D. ROLLBACK

Answer: C

Explanation:

When a user's Server Process writes a sort chunk to disk, it writes the data to the user's temporary tablespace. This tablespace, although it is referred to as the user's temporary tablespace, can have the tablespace attribute of being either permanent or temporary. If there is no TEMPORARY TABLESPACE specified for the user, SYSTEM tablespace will be considered as temporary tablespace for the user.

Incorrect Answers

A: There is no TEMP tablespace in Oracle by default.

B: There is no DATA tablespace in Oracle by default.

D: ROLLBACK tablespace cannot be used as temporary tablespace.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 420-423

Chapter 8: Tuning Disk I/O

QUESTION 15:

Which dynamic view is most useful for determining the current number of blocks allocated to a buffer pool?

- A. V\$CACHE
- B. V\$SESS_IO
- C. V\$SYSSTAT
- D. V\$BUFFER_POOL

Answer: D

Explanation:

The V\$BUFFER_POOL dynamic performance view contains information about the configuration of the multiple Buffer Pools themselves. You can monitor the performance of the Buffer Pools using the V\$BUFFER_POOL and V\$BUFFER_POOL_STATISTICS dynamic performance views.

Incorrect Answers

A: V\$CACHE dynamic view is an Oracle9i Real Application Clusters view. This view contains information from the block header of each block in the SGA of the current instance as related to particular database objects.

B: This view lists I/O statistics for each user session.

C: Sort activity can be monitored using the V\$SYSSTAT and V\$SORT_SEGMENT

dynamic data dictionary views, using the output from STATPACK and REPORT.TXT, and using the output from the OEM Performance Manager.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 274-275

Chapter 5: Tuning the Database Buffer Cache

QUESTION 16:

Which three statements about improving the performance of the database buffer cache by creating multiple buffer pools are true? (Choose three)

- A. One, two, or three pools may be defined.
- B. There are at least 50 blocks per LRU latch for each pool.
- C. Each buffer pool is assigned latches taken from DB_BLOCK_LRU_LATCHES.
- D. The size of the DEFAULT pool is obtained by adding all the pools to the value of the DB_BLOCK_BUFFERS parameter.

Answer: A, B, C

Explanation:

Unless you specify otherwise in the init.ora, only the Default Pool is created at instance startup. But Oracle provides you also with the ability to divide the Database Buffer Cache into as many as three separate areas called Buffer Pools: Keep Pool (optional), Recycle Pool (optional) and Default Pool (mandatory). There are at least 50 blocks per LRU latch for each pool. Also each buffer pool is assigned latches taken from DB_BLOCK_LRU_LATCHES.

Incorrect Answers

D: Default Pool is used to cache segments that are not designated for either the Keep or Recycle pools. The size of this pool is designated in bytes, kilobytes, megabytes, or gigabytes, by the init.ora parameter DB_CACHE_SIZE. Unlike Oracle8i, where the memory for the Keep and Recycle pools was taken from the memory allocated to the Default Pool, Oracle9i independently assigns the memory to each of the three Buffer Pool types.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 265-266, 468-476

Chapter 5: Tuning the Database Buffer Cache, Chapter 9: Tuning Contention

QUESTION 17:

Which V\$ view which is used to monitor granule allocation and de-allocation in the buffer cache?

- A. V\$CACHE
- B. V\$SESSTAT
- C. V\$BUFFER_POOL
- D. V\$LIBRARYCACHE

Answer: C

Explanation:

This view displays information about all buffer pools available for the instance, including granule allocation and de-allocation in the buffer cache.

Incorrect Answers

A: This is an Oracle9i Real Application Clusters view. This view contains information from the block header of each block in the SGA of the current instance as related to particular database objects.

B: This view lists user session statistics.

D: This view contains statistics about library cache performance and activity.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 275-277
Chapter 5: Tuning the Database Buffer Cache

QUESTION 18:

What will this statement do?

```
CREATE TABLESPACE temp  
DATAFILE 'C:\database\temp.dbf' SIZE 10n  
Temporary;
```

- A. Create a tablespace that will be dropped on instance shutdown.
- B. Create a tablespace in which the user can create segments for usage during sorts.
- C. Create a tablespace in which Oracle can create segments for usage during sorts.
- D. Create a tablespace in which a user can create tables that will be automatically dropped after a week.

Answer: C

Explanation:

This command creates a tablespace in which Oracle can create segments for usage during sorts.

Incorrect Answers

A: You cannot drop the default temporary tablespace until another has been created because doing so would leave the database with nowhere to perform to-disk sorts. Unlike a sort segment stored in a permanent tablespace, the sort segment in the temporary tablespace is not dropped when the user's sort completes. Instead, the first sort operation following instance startup creates a sort segment that remains in the temporary tablespace for reuse by subsequent users who also perform sorts to disk. This sort segment will remain in the temporary tablespace until instance shutdown. So only a sort segment will be dropped, not entire tablespace on instance shutdown.

B: User itself cannot create any objects in the temporary tablespace: it is used exclusively by Oracle.

D: User itself cannot create any objects in the temporary tablespace: it is used exclusively by Oracle.

QUESTION 19:

Which type of transaction should you assign to a specific large rollback segment?

- A. Batch jobs that modify many rows.
- B. Long running serializable transactions.
- C. Long running reports, to avoid 'snapshot too old' errors.
- D. Discrete transactions that modify many rows in the same block.

Answer: A

Explanation:

You need to assign a batch jobs modifying many rows to specific large rollback segment using SET TRANSACTION command.

Incorrect Answers

- B: Long running serializable transactions do not require large rollback segments.
- C: Reports do not modify any table records. So they do not need specific rollback segment.
- D: Discrete transactions that modify many rows in the same block will not cause rollback segment contention.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 187-188
Chapter 3: SQL Application Tuning and Design

QUESTION 20:

What is the least number of buffers an LRU latch must cover in the database buffer cache?

- A. 5
- B. 10
- C. 30
- D. 50
- E. 100

Answer: D

Explanation:

There are at least 50 blocks per LRU latch for each pool.

Incorrect Answers

- A: There are at least 50, not 5, blocks per LRU latch for each pool.
- B: There are at least 50, not 10, blocks per LRU latch for each pool.
- C: There are at least 50, not 30, blocks per LRU latch for each pool.
- E: There are at least 50, not 100, blocks per LRU latch for each pool.

QUESTION 21:

Which three statements about rebuilding indexes are true? (Choose three)

- A. The ALTER INDEX REBUILD command is used to change the storage characteristics of an index.
- B. Using the ALTER INDEX REBUILD is usually faster than dropping and recreating an index because it uses the fast full scan feature.
- C. Oracle8i allows for the creation of an index or re-creation of an existing index while allowing concurrent operations on the base table.
- D. When building an index, the NOLOGGING and UNRECOVERABLE keywords can be used concurrently to reduce the time it takes to rebuild.

Answer: A, B, C

Explanation:

You can use the ALTER INDEX REBUILD command to optimize the storage characteristics of an index.

Using the ALTER INDEX command with the REBUILD option is an effective way to quickly rebuild an index because the existing index entries are used to create the new index. The ONLINE option of the ALTER INDEX REBUILD command should be used to minimize any locking issues that occur when an index is rebuilt while users continue to perform DML on the index's underlying table.

Incorrect Answers

D: When building an index, the NOLOGGING and UNRECOVERABLE keywords cannot be used concurrently to reduce the time it takes to rebuild.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 151-152
Chapter 3: SQL Application Tuning and Design

QUESTION 22:

Where can you find the nondefault parameters when the instance is started?

- A. Alert log
- B. Online redo log
- C. Archiver redo log
- D. SYSTEM user's trace file

Answer: A

Explanation:

Alert log file shows the nondefault parameters when the instance is started.

Incorrect Answers

- B: Online redo log is not used to store this information.
 - C: Archive redo log is just archived copy of online redo log. It does not have any information about nondefault parameters.
 - D: SYSTEM's user's trace file is not used to provide this information.
- OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 34-38
Chapter 2: Sources of Tuning Information
-

QUESTION 23:

What should be two goals in tuning rollback segments? (Choose two)

- A. Transactions should never wait for access to rollback segment.
- B. No transaction, however large or exceptional, should ever run out of rollback space.
- C. Rollback segments should be configured to extend continually during normal processing.
- D. The ratio of waits to the rollback segment header blocks should be less than 5% of the sum of access.

Answer: A, B

Explanation:

There are two main goals for rollback segments: transactions should never wait for access to rollback segment; no transaction should ever run out of rollback space because it used to keep the read-consistent view of the changed data.

Incorrect Answers

- C: Rollback should not extend continually during normal processing. It is possible only as exception to keep data for batch jobs performing DML operations with many rows.
- D: Transactions should never wait for access to rollback segment.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 425-429
Chapter 8: Tuning Disk I/O

QUESTION 24:

Which statement about improving the performance of the database buffer cache by creating multiple buffer pools is true?

- A. The KEEP buffer pool must also be deferred if the RECYCLE pool is defined.
- B. The buffer pool for an object can be set explicitly only at object creation time.
- C. The blocks from an object without an explicitly set buffer pool go into the RECYCLE pool.
- D. Buffer pools are assigned to a segment, so option with multiple segments can have blocks in multiple buffer pools.

Answer: D

Explanation:

Oracle provides you with the ability to divide the Database Buffer Cache into as many as three separate areas called Buffer Pools. Segments are then explicitly assigned to use the appropriate Buffer Pool as determined by the DB

A. Option with multiple segments can have blocks in multiple buffer pools.

Incorrect Answers

A: There is no such relation with the KEEP buffer pool and the RECYCLE buffer pool: they can be defined independently.

B: The buffer pool for an object can be changed after an object creation.

C: The blocks from an object without an explicitly set buffer pool go into the DEFAULT pool.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 265-274
Chapter 5: Tuning the Database Buffer Cache

QUESTION 25:

What should one be your tuning goals?

A. Use as much memory as possible.

B. Use multiple copies of the code in memory.

C. Access the most possible number of blocks from disk.

D. Access the least possible number of blocks from disk.

Answer: D

Explanation:

The main goal is to access the least possible number of blocks from disk because I/O operations are significantly more expensive as memory operations.

Incorrect Answers

A: You need to use as less memory as possible.

B: You need to share the same code in the memory to reduce the memory usage.

C: The main goals is to access the least, not the most, possible number of blocks from disk

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 6-10
Chapter 1: Introduction to Performance Tuning

QUESTION 26:

When should you recommend changing the application in order to reuse more SQL?

A. When the GETHITRATIO in the V\$LIBRARYCACHE view is above 0.99.

B. When the misses in the dictionary cache are greater than 1% of the hits.

C. When the ratio of GETHITS to GETS in the V\$LIBRARYCACHE view is less then 0.9.

D. When the ratio of RELOADS to PINS in the V\$LIBRARYCACHE view is less than

0.01.

Answer: C

Explanation:

The ratio of parsed statements (GETS) to those that did not require parsing (GETHITS) is calculated in the GETHITRATIO column of V\$LIBRARYCACHE. The higher this number is, the better the application is performing.

Incorrect Answers

A: When the GETHITRATIO in the V\$LIBRARYCACHE view is above 0.99, application performance is good.

B: The dictionary cache has nothing to do with SQL statements: it stores the data dictionary information in the memory.

D: The RELOADS column in the V\$LIBRARYCACHE view shows the number of times that an executed statement had to be re-parsed because the Library Cache had aged out or invalidated the parsed version of the statement. Reload activity can be monitored by comparing the number of statements that have been executed (PINS) to the number of those statements that required a reload (RELOADS). The less this number is, the better the application is performing.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 210-213
Chapter 4: Tuning the Shared Pool

QUESTION 27:

What are two possible causes of lock contention? (Choose two)

- A. Uncommitted changes.
- B. Too many rollback segments.
- C. Improperly sized redo logs.
- D. Shared pool is sized too large.
- E. Other protocols imposing unnecessarily high locking levels.

Answer: A, E

Explanation:

Lock contention can arise because of uncommitted changes and unnecessarily table level locks.

Incorrect Answers

B: Large amount of rollback segments decrease possibility of lock contention.

C: Lock contention is not related with the size of redo logs.

D: Lock contention is not related with shared pool size.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 480-487
Chapter 9: Tuning the Contention

QUESTION 28:

Which component will NEVER allocate memory from the large pool?

- A. Oracle Library Cache.
- B. Oracle Parallel Query.
- C. Oracle Recovery Manager.
- D. Oracle Multithreaded Server.

Answer: A

Explanation:

Oracle Library Cache will NEVER allocate memory from the Large Pool.

Incorrect Answers

B: Oracle Parallel Query (PQ) can cause the Shared Pool to cache additional items not related to application SQL. That's why it can allocate memory from the Large Pool.

C: Oracle Recovery Manager (RMAN) utility can use the Large Pool to process requests for optional features.

D: UGA information for the Shared Server option is cached in the Large Pool if it's defined. Otherwise, it uses the Shared Pool.

OCP: Oracle 9i Performance Tuning Study Guide, Jaseph C. Johnson, p. 308-310

Chapter 8: Tuning Other SGA Areas

QUESTION 29:

Database Resource Manager uses resource plans to determine resource limits for the set of users. Which statement is true in reference to resource plans?

90

- A. Resource plans are set using profiles.
- B. Only one resource plan can be stored in the database at one time.
- C. The database can have many resources plans, but only one can be active at any one time.
- D. The database can have many resources plans, and each user chooses which plan to belong to.

Answer: C

Explanation:

Oracle8i introduced a new feature, Resource Manager, which was designed to improve the allocation and management of server resources needed by application users. In Oracle 9i, it is possible to control numerous aspects of application processing via Resource Manager. The database can have many resources plans, but only one can be active at any one time.

Incorrect Answers

A: Resource plans are not set using profiles.

B: The database can have many resources plans.

D: Each user cannot choose which plan to belong to because only one plan can be active

at any one time.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 517-547
Chapter 10: Operating System Tuning

QUESTION 30:

Which three actions will cause queries to place a table's blocks at the most-recently-used end of the LRU list for the DEFAULT pool? (Choose three)

- A. Creating a table with the CACHE option.
- B. Querying the table by using a CACHE hint.
- C. Ensuring the query performs a full table scan.
- D. Defining the table without the option for caching.
- E. Altering an existing table to set the CACHE option.
- F. Ensuring the query does not retrieve data through index lookup.
- G. Creating a separate database buffer cache to hold cached table.

Answer: A, B, E

Explanation:

If you are creating a table with the CACHE option or altering an existing table to set the CACHE option will place a table's blocks at the most-recently-used end of the LRU list. Querying the table by using a CACHE hint will cause queries to place a table's blocks at the most-recently-used end of the LRU list.

Incorrect Answers

- C: It does not mean that blocks will be placed at the most-recently-used end of the LRU list if the query performs a full table scan.
 - D: Defining the table without the option for caching will not place blocks in the LRU list.
 - F: This statement will not cause queries to place a table's blocks at the most-recently-used end of the LRU list.
 - G: You don't need to create a separate database buffer cache to hold cached table.
- OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 276-279
Chapter 5: Tuning the Database Buffer Cache
-

QUESTION 31:

What is the main reason to create a reverse key index on a column?

- A. The column is populated using a sequence.
- B. The column contains many different values.
- C. The column is mainly used for value range scans.
- D. The column implementing an inverted list attribute.

Answer: A

Explanation:

The Reverse Key Index (RKI) is a special type of B-Tree index. The RKI is useful when an index is built on a column that contains sequential numbers.

Incorrect Answers

B: B-Tree index will be appropriate choice for the column with many different values.

C: The column is mainly used for value range scans is not good candidate for the Reverse Key Index. Reverse Key indexes are only useful for equality and non-equality searches. Queries that perform range scans (e.g., using BETWEEN, >, <) on columns that are Reverse Key indexed will not be able to use the index and will cause full table scans.

D: The RKI does not work for the column implementing an inverted list attribute.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 156-158
Chapter 3: SQL Application Tuning and Design

QUESTION 32:

Which type of table is the best candidate to be cached?

- A. Small table rarely retrieved with a full table scan.
- B. Large table rarely retrieved with a full table scan.
- C. Small table frequently retrieved with a full table scan.
- D. Large table frequently retrieved with a full table scan.

Answer: C

Explanation:

Small tables frequently retrieved with a full table scan can be cached.

Incorrect Answers

A: If table is rarely retrieved there is no sense to cache it.

B: Large table cannot be easy cached because CACHE option works mostly for the small tables that can be loaded in the memory. Also if table is rarely retrieved there is no sense to cache it.

D: Large table is not good candidate to be cached.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 276-279
Chapter 5: Tuning the Database Buffer Cache

QUESTION 33:

Which initialization parameter specifies the location of the alert log file?

- A. UTL_FILE_DIR
- B. USER_DUMP_DEST
- C. LOG_ARCHIVE_DEST
- D. BACKGROUND_DUMP_DEST

Answer: D

Explanation:

BACKGROUND_DUMP_DEST initialization parameter specifies the location of the alert log file.

Incorrect Answers

A: UTL_FILE_DIR lets you specify one or more directories that Oracle should use for PL/SQL file I/O.

B: USER_DUMP_DEST is used to specify the location of user trace files.

C: LOG_ARCHIVE_DEST shows the directory for the archived redo logs.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 40-41
Chapter 2: Sources of Tuning Information

QUESTION 34:

The NOLOGGING mode in SQL statements is a tool used to reduce redo operations, but NOLOGGING does not apply to every operation for which the attribute is set. Which three SQL statements can use the NOLOGGING mode to reduce redo operations?

(Choose three)

- A. UPDATE
- B. CREATE INDEX
- C. ALTER INDEX...REBUILD
- D. Conventional Path INSERT
- E. CREATE TABLE.... AS SELECT

Answer: B, C, E

Explanation:

You can use the NOLOGGING mode to create index, to rebuild index and to create table as select from the other table.

Incorrect Answers

A: It cannot be used for DML operations itself. Only table can be switched to this mode and than UPDATE command can be performed for the table.

D: Conventional Path INSERT cannot use it. Once the NOLOGGING attribute is set on a table, redo entry generation will be suppressed for all subsequent DML on the table only when that DML is of the following types: Direct Path loads using SQL*Loader, direct load inserts using the /*+ APPEND*/ hint.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 335-336
Chapter 7: Tuning Redo Mechanisms

QUESTION 35:

Which two statements about database blocks are true? (Choose two)

- A. DSS environment prefer a large block size.
- B. Small block sizes result in more block contention.
- C. Random access to large object favours a large block size.
- D. You can reduce the number of block visits by packing rows as closely as possible into

blocks.

E. To change the database block size, you must shut down the instance and perform a STARTUP RESETLOGS after you make the change.

Answer: A, D

Explanation:

DSS database will work better with a large database block size due to lot of full scans. Users of these systems are concerned with response time, which is the time it takes to get the results from their queries, so data need to be packed as closely as possible into blocks.

Incorrect Answers

B: Small block sizes will cause overload in the DSS system. OLTP systems works better with small block sizes.

C: Random access to large object does not prefer a large block size.

E: You cannot change a database block size with the STARTUP RESETLOGS command. There is no STARTUP RESETLOGS command in Oracle. The STARTUP and ALTER DATABASE OPEN RESETLOGS commands exist.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 187-188
Chapter 3: SQL Application Tuning and Design

QUESTION 36:

The ORDERS table has millions of rows and is accessed very often with an index (ORDID_NDX) on a primary key (ORD_ID). Where should ORDERS and ORDID_NDX be stores?

- A. Same tablespace
- B. Different tablespace on the same disk.
- C. Tablespace containing a rollback segment.
- D. Different tablespaces on different disks.

Answer: D

Explanation:

To avoid I/O contention you need to store different tablespaces on different disks.

Incorrect Answers

A: Storing data and indexes on the same tablespace can cause performance degradation.

B: Different tablespace on the same disk will not fix I/O contention problem.

C: Rollback tablespace should be stored way far from the data and indexes.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 368-388
Chapter 8: Tuning Disk I/O

QUESTION 37:

Which two statements about row migration are true? (Choose two)

1Z0-033

- A. Row migration is caused by a PCTREE value set too low.
- B. Row migration can be resolved using the ANALYZE command.
- C. Row migration can be reduced by choosing a larger block size.
- D. Row migration means that row pieces are stored in different blocks.
- E. Queries that use an index to select migrated rows perform additional I/O.

Answer: A, B

Explanation:

Row migration occurs when a previously inserted row is updated. If the update to the row causes the row to grow larger than the space available in the block specified by PCTFREE, Oracle moves (or migrates) the row to a new block. The ANALYZE command just populates the CHAIN_CNT column of DBA_TABLES, which is otherwise null. This column will indicate how many of a table's rows are using more than one block to store data. However, no distinction is made between the number of chained rows and the number of migrated rows.

Incorrect Answers

C: Row migration cannot be reduced by choosing a larger block size. Rebuilding the database with a larger block size can give you a significant performance benefit because of row chaining, but not row migration problem elimination.

D: Row migration is migration the row to a new block. Whenever a row split over two or more multiple blocks it is referred to as a chained row.

E: Queries that use an index to select migrated rows do not perform additional I/O.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 405-407
Chapter 8: Tuning Disk I/O

QUESTION 38:

What are three indications of contention for this rollback segment header? (Choose three)

- A. A nonzero value in the WAITS column of the V\$ROLLSTAT view.
- B. A nonzero value in the UNDO HEADER column of the V\$WAITSTAT view.
- C. A nonzero value in the ROLL_SEG_WAIT column of the V\$ROLLSEGS view.
- D. A nonzero value in the UNDO_HEADER_WAITS columns of the V\$ROLLBACK_SEGS view.
- E. A nonzero value in the Undo Segment TX Slot event of the V\$SYSTEM_EVENT view.

Answer: A, B, E

Explanation:

The ROLLSTAT view contains detailed information regarding the behavior of the rollback segments in the database. In particular, the columns USN, GETS and WAITS are particularly useful for measuring contention for the rollback segment's header. WAITS column shows number of times a user Server Process needed to access the

rollback segment header and experienced a wait. The V\$WAITSTAT view contains information on block contention statistics. A nonzero value in the UNDO_HEADER column indicates a contention problem. The V\$SYSTEM_EVENT view tracks performance related information on rollback segments via the Undo Segment TX Slot statistic. Ideally, the value in the Undo Segment TX Slot event should be consistently at or near zero.

Incorrect Answers

C: There is no V\$ROLLSEGS data dictionary view in the Oracle.

D: There is no V\$ROLLBACK_SEGS data dictionary view in the Oracle.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 425-441

Chapter 8: Tuning Disk I/O

QUESTION 39:

When tables are stored in locally managed tablespaces, where is extent allocation information stored?

- A. Memory
- B. Data dictionary.
- C. Temporary tablespace.
- D. Within the locally managed tablespace

Answer: D

Explanation:

The extent allocation information is stored in the corresponding tablespace if tables are stored in locally managed tablespaces.

Incorrect Answers

A: This information is stored in the corresponding tablespace, not in the memory.

B: Only for tables in data dictionary managed tablespaces extent allocation is stored in the data dictionary.

C: Temporary tablespaces are not used to store extent allocation information.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 377

Chapter 8: Tuning Disk I/O

QUESTION 40:

What is one difference between I/O slaves and DBWn processes for the DB Writer?

- A. In Oracle81, I/O slaves are not available; only DBWn processes are available
- B. I/O slaves perform the write function only, while DBWn processes also perform data-gathering activity.
- C. I/O slaves will work only with synchronous I/O, whereas DBWn processes are available only within asynchronous I/O.
- D. I/O slaves will work only with asynchronous I/O, whereas DBWn processes are available only within synchronous I/O.

Answer: B

Explanation:

Database Writer slave processes are similar to the actual DBW0 process itself, except they can only perform write operations, not move buffers from the LRU List to the Dirty List in the Database Buffer Cache as DBW0 does. The purpose of these slaves is to simulate asynchronous I/O on systems that only support synchronous I/O.

Incorrect Answers

A: I/O slaves are available in the Oracle8i.

C: I/O slaves will work not only with synchronous I/O, they can simulate asynchronous I/O operations.

D: I/O slaves work with synchronous I/O and it can simulate asynchronous I/O operations.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 393-441

Chapter 8: Tuning Disk I/O

QUESTION 41:

With reference to Oracle data storage structures, a cluster is defined as?

A. A group of table that each have more then 2 low cardinality columns.

B. A data structure where a group of one or more tables have their own dedicated tablespaces.

C. A group of one or more tables which resides in a tablespace that is striped across multiple disks.

D. A group of one or more tables that share the same data blocks because they share common columns and are often used together in join queries.

Answer: D

Explanation:

Clusters are used to store the data from one or more tables in the same physical Oracle blocks. In general, clustered tables should have these attributes: always be queried together and only infrequently on their own, have little or no DML activity performed on them after the initial load, have roughly equal numbers of child records for each parent key.

Incorrect Answers

A: Cluster definition is not related with columns' cardinality.

B: Clusters' tables do not have their own dedicated tablespaces. They use the same physical Oracle blocks.

C: Tablespace can be stripped across multiple disks (via its datafiles), but tables in side this tablespace can or cannot be a cluster.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 182-186

Chapter 3: SQL Application Tuning and Design

QUESTION 42:

You have a table with a million rows. You want to build an index on a column in the table that has a low cardinality. The table is part of a Decision Support System. Your goal is to build an index that would be efficient for queries using AND/OR predicates. Which type of index would be most suitable?

- A. B-Tree Index.
- B. Bitmap Index.
- C. Reverse Key Index.
- D. Compresses Indexes.

Answer: B

Explanation:

Because of low cardinality of column in the table of DSS system bitmap index will be useful only to perform queries using AND/OR predicates.

Incorrect Answers

A: B-tree index works better for the OLTP systems with high cardinality columns.

C: Reverse Key Index is useful when an index is built on a column that contains sequential numbers.

D: There is no compressed index type in Oracle.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 187-188
Chapter 3: SQL Application Tuning and Design

QUESTION 43:

What are two main benefits of index-organized tables? (Choose two)

- A. More concurrency.
- B. Faster full table scans.
- C. Fast primary key-based access.
- D. Less contention on the segment header.
- E. Less storage is required because there is no duplication of primary key values.

Answer: C, E

Explanation:

Oracle provides Index Organized Tables (IOTs) to store a table's data in a specific order. Instead of storing a row ID pointer to where the rest of the row data is stored, the row data is actually stored in its entirety in the index itself. The table rows are stored in index order. If you access the table using its primary key, an IOT will return the rows more quickly than a traditional table. The extra free space is available because there is no duplication of primary key values in an IOT.

Incorrect Answers

A: IOTs do not provide more concurrency.

B: They will not be helpful for full table scans.

D: Less contention on the segment header is not a main benefit of index-organized tables.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 158-162

Chapter 3: SQL Application Tuning and Design

QUESTION 44:

Which four statements are true regarding materialized views? (Choose four)

A. Materialized views cannot be partitioned, nor can they be defined on partitioned tables.

B. Materialized views are often used in data warehouses to increase the speed of queries on very large datatables.

C. Queries that benefit from the use of materialized views often involve joins between tables or aggregations such as SUM.

D. A materialized view stores both the definition of a view and the rows resulting from the execution of the views.

E. Materialized views can be used to replicate data, which was formerly achieved using the CREATE SNAPSHOT statement.

Answer: B, C, D, E

Explanation:

Stored outlines help speed up queries by telling the optimizer how to tackle the query execution associated with a particular SQL statement. Materialized views are also designed to speed up queries by storing data from queries in a pre-joined, pre-summarized format. Unlike a traditional view, which is merely stored in the data dictionary as a SELECT statement that is executed when the view is accessed, a materialized view stores the physical results of the view in its own segment, separate and distinct from the underlying table on which the view is based. Materialized views are intended primarily for use in data warehouses and Decision Support Systems where large volumes of data are accessed and summarized using queries.

Incorrect Answers

A: Materialized view segment can be stored in its own tablespace and can be indexed and partitioned.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 140-146

Chapter 3: SQL Application Tuning and Design

QUESTION 45:

Which statement is valid regarding index clusters?

A. Index clusters can only be used for tables with low cardinality columns.

B. Index clusters are generally well suited for tables that have many full table scans.

C. Normal B-Tree indexes do not store null key values, whereas cluster indexes store null keys.

D. A cluster index always takes up much more storage space than a normal index for the same set of key values.

Answer: C

Explanation:

Index clusters can store null keys. B-Tree indexes do not.

Incorrect Answers

A: Index clusters cannot only be used for tables with low cardinality columns. They are used to store the data from one or more tables in the same physical Oracle blocks.

B: They are not very good suited for tables with many full table scans.

D: A cluster index does not use much more storage space than a normal index.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 182-186

Chapter 3: SQL Application Tuning and Design

QUESTION 46:

You have a table called COMPANY created with the following SQL in your database:

```
CREATE TABLE company (  
company_id NUMBER(5) not null,  
company_name VARCHAR2(4) NOT NULL  
)  
NOLOGGING  
PCTFREE 0  
storage (initial 32K next 32K pctincrease 0)
```

You have created 2 indexes, one on the COMPANY_ID column and the other on the COMPANY_NAME column. Evaluate these 4 SQL statements, assuming use of the Rule-Bases Optimizer:

- 1.SELECT*FROM company Where company_id=40;
- 2.SELECT*FROM company Where company_id+0=40;
- 3.SELECT*FROM company Where company_name='oeacle';
- 4.SELECT*FROM company Where company_name=||"oeacle";.

What is a valid conclusion about index usage in the above 4 SQL statements?

- A. All 4 SQL statements will use an index.
- B. Statements 1, 2 & 3 will use an index, and in statement 4 the index will be ignored.
- C. Statements 1, 3 & 4 will use an index, and in statement 2 the index will be ignored.
- D. None of the SQL statements will use an index.
- E. Statements 1 & 3 will use an index, and in statement 2 & 4 the index will be ignored.
- F. Only statement 1 will use an index, and in statement 2, 3 & 4 the index will be ignored.

Answer: E

Explanation:

Statements 1 & 3 will use an index, and in statement 2 & 4 the index will be ignored.

Incorrect Answers

A: In statement 2 & 4 the index will be ignored.

B: In statement 2 the index will be ignored.

C: In statement 4 the index will be ignored.

D: Statements 1 & 3 will use an index.

F: Statements 1 & 3 will use an index.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 109

Chapter 3: SQL Application Tuning and Design

QUESTION 47:

What is the main reason for a row overflow area when creating index-organized tables?

A. To avoid row chaining and migration.

B. To speed up full table scans and fast full index scans.

C. To improve performance when the index-organized table is clustered.

D. To keep the B-Tree structure densely clustered to allow more rows per leaf block.

Answer: D

Explanation:

The main reason to use a row overflow for IOT is the need to keep the B-Tree structure densely clustered to allow more rows per leaf block. During IOT table creation you can specify the OVERFLOW tablespace where the second half of the row data will be stored when the row's length exceeds the size set aside in PCTTHRESHOLD.

Incorrect Answers

A: The row overflow will not help you to avoid row chaining because it's used exactly to keep chained rows for better balance of the B-Tree structure.

B: It will not speed up full table scans and fast full index scans because it will be required additional time to read data from the overflow area.

C: The row overflow area will not improve an overall performance of IOT.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 158-162

Chapter 3: SQL Application Tuning and Design

QUESTION 48:

Which two statements correctly describe the use of the REFRESH option when creating materialized views? (Choose two)

A. Use the REFRESH COMPLETE option to truncate the materialized view and repopulate the view with data from the base tables in the query.

B. Use the REFRESH FAST option to populate the materialized view data from the base tables that has changed since the last re-sync.

C. Use the REFRESH FAST option to truncate the materialized view and populate the view with data from the base tables in the query.

1Z0-033

- D. Use the REFRESH FAST ON DEMAND option to repopulate the materialized view with data from base tables after each commit to any of the base tables.
- E. Use the REFRESH COMPLETE option to update the existing data in the view with all the new changes from the base tables since the last re-sync, without truncating the existing materialized view.

Answer: A, B

Explanation:

Materialized views are designed to speed up queries by storing data from queries in a pre-joined, pre-summarized format. If you do want the data to be kept in sync, you will specify either the COMPLETE, FAST, or FORCE option during the creation of the view. If you use a COMPLETE option, during a refresh, the materialized view is truncated and then completely repopulated with data from the base tables in the query. Using the FAST option, the materialized view is populated only with data that has changed in the base table since the last re-sync. This refresh is performed using the view's log data or by ROWID.

Incorrect Answers

C: The REFRESH FAST option does not truncate the materialized view.

D: ON DEMAND option defines that you can also manually refresh the contents of a materialized view.

E: Using the REFRESH COMPLETE option the materialized view is truncated and then completely repopulated with data from the base tables

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 140-146

Chapter 3: SQL Application Tuning and Design

QUESTION 49:

The DBA's task of building a well performing database often begins with selecting proper data storage structures. DBAs should be aware of what types of storage structures are appropriate for various data access methods.

Which three data access methods will enhance database performance when combined with the appropriate types of application? (Choose three)

- A. Cluster
- B. Advanced Queue
- C. Materialized view
- D. Advanced Replication
- E. Index-organized table
- F. Real Application Cluster

Answer: A, C, E

Explanation:

Clusters, materialized views and index-organized tables will enhance database performance if they are used with appropriate types of applications. For example,

Decision Support Systems (DSS) and data warehouse make heavy use of full table scans so the appropriate use of indexes and hash clusters are important. Index-organized tables can be also important tuning options for large DSS systems.

Incorrect Answers

B: Advanced Queue will not enhance database performance.

D: Advanced Replication is not a data access method.

F: Real Application Cluster is not a data access method, it's just data storage structure.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 140-146, 158-162, 182-186

Chapter 3: SQL Application Tuning and Design

QUESTION 50:

What are two main advantages of using bitmap indexes? (Choose two)

A. Bitmap indexes use less storage space.

B. Bitmap indexes offer maximum concurrency.

C. Bitmap indexes are easy to maintain when you issue DML statements.

D. Bitmap segments are updated upon COMMIT, at the end of the transaction.

E. Bitmap indexes work very fast with multiple predicates that are combined with AND, OR, and NOT operators.

Answer: A, E

Explanation:

Bitmap indexes use less storage space as B-Tree indexes do. Unlike B-Tree indexes, bitmap indexes create a binary mapping of the rows in the table and store that map in the index blocks. This means the resulting index will require significantly less space to store the index data and retrieve the rows of an equality match on the indexed column more quickly than an equivalent B-Tree index. They work very fast with multiple predicates also.

Incorrect Answers

B: They do not offer maximum concurrency. Regular B-Tree indexes can provide it.

C: Bitmap indexes should not be used on tables that have high INSERT, UPDATE, or DELETE activity. These DML operations are costly in terms of performance because they cause locking to occur at the bitmap level and require that the entire bitmap for all possible values be rebuilt dynamically.

D: Bitmap segments are updated before COMMIT, not waiting the end of the transaction.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 153-155

Chapter 3: SQL Application Tuning and Design

QUESTION 51:

In an index-organized table, what type of segment is used to store row data that exceeds the index's PCTTHRESHOLD?

- A. DATA segment.
- B. INDEX segment.
- C. CHAIN segment.
- D. EXCESS segment
- E. OVERFLOW segment.

Answer: E

Explanation:

The main reason to use a row overflow for IOT is the need to keep the B-Tree structure densely clustered to allow more rows per leaf block. During IOT table creation you can specify the OVERFLOW tablespace where the second half of the row data will be stored when the row's length exceeds the size set aside in PCTTHRESHOLD.

Incorrect Answers

A: DATA tablespace is not used to store row data that exceeds the index's PCTTHRESHOLD. It's used for regular data.

B: INDEX tablespace is not used to store row data that exceeds the index's PCTTHRESHOLD. It's used to store indexes.

C: There is no CHAIN segment exist for the IOT.

D: There is no EXCESS segment exist for the IOT.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 158-162
Chapter 3: SQL Application Tuning and Design

QUESTION 52:

The optimizer rewrites a query so that the query can access a materialized view instead of the base tables. Although query rewrite activity is transparent to the applications, there are certain prerequisites that need to be satisfied for the optimizer to rewrite queries. Which statement correctly describe one of the prerequisites?

- A. OPTIMIZER_MODE must be set with cost-based optimization.
- B. QUERY_REWRITE_ENABLED must be specified in the parameter file.
- C. A user who owns the materialized view must also own PLAN_TABLE.
- D. A user must be granted QUERY REWRITE system privilege to enable materialized views in any schema.
- E. QUERY_REWRITE_ENABLED=TRUE must be included in the option clause when a materialized view is created.

Answer: A

Explanation:

To rewrite queries for better performance OPTIMIZER_MODE should be set with cost-based optimization.

Unlike the rule-based optimizer that has only its predefined guidelines to follow when executing a query, the cost-based optimizer (CBO) considers many different execution plans and then selects the one with the lowest execution plan. The

QUERY_REWRITE_ENABLED init.ora parameter allows optimizer to dynamically rewrite queries to take advantage of materialized views when set to TRUE.
QUERY_REWRITE_INTEGRITY determines the degree to which the data consistency is to be adhered to when accessing materialized views. This will work only if OPTIMIZER_MODE is set to one of the cost-based optimizer modes.

Incorrect Answers

B: QUERY_REWRITE_ENABLED can be specified in the parameter file, but there is not requirement.

C: There is no prerequisite that a user who owns the materialized view must also own PLAN_TABLE.

D: It is not required that a user must be granted QUERY REWRITE system privilege to enable materialized views in any schema.

E: QUERY_REWRITE_ENABLED=TRUE may be set only in the init.ora file to allow optimizer to dynamically rewrite queries to take advantage of materialized views. It cannot be used in the option when a materialized view is created.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 140-146
Chapter 3: SQL Application Tuning and Design

QUESTION 53:

Which two statements are true with respect to hash clusters? (Choose two)

- A. Hash clusters perform well when the cluster keys are updated rarely.
- B. Hash clusters perform well when the cluster keys are updates frequently.
- C. Use of hash clusters may be beneficial for a data warehouse type of application.
- D. Full table scans are generally faster on clustered tables than on non-clustered tables.
- E. If an application mostly issues range searches, hash key is usually found in a single read while a nonclustered table with an index requires a minimum pf 2 I/O's

Answer: A, C

Explanation:

Hash clusters are used in place of a traditional index to quickly find rows stored in a table. Like cluster indexes, the usefulness of Hash clusters is also limited to very specific situations. They can improve performance when there is little or no DML activity performed on them (especially with cluster keys) after the initial load. They can be very helpful in the DSS or data warehouse systems with many heavy usage of table scans.

Incorrect Answers

B: Hash clusters perform poorly when the cluster keys are updates frequently.

D: Full table scans are generally faster on non-clustered tables because they do not require to calculate a hash function to retrieve each row.

E: This statement about number of I/O operations will work only for one value search, not for range searches. In case of one value search, really only one I/O operation will be required to retrieve the row from the hash table, instead of the two or more with the relational table/B-Tree index combination.

QUESTION 54:

What is the main reason to create a reverse key index on a column?

- A. The column contains many different values.
- B. The column is mainly used for value range scans.
- C. The column is populated using sequential numbers.
- D. The column implements an inverted list attribute.

Answer: C

The Reverse Key Index (RKI) is a special type of B-Tree index. The RKI is useful when an index is built on a column that contains sequential numbers.

Explanation:

Incorrect Answers

- A: B-Tree index will be appropriate choice for the column with many different values.
 - B: The column is mainly used for value range scans is not good candidate for the Reverse Key Index. Reverse Key indexes are only useful for equality and non-equality searches. Queries that perform range scans (e.g., using BETWEEN, >, <) on columns that are Reverse Key indexed will not be able to use the index and will cause full table scans.
 - D: The RKI does not work for the column implementing an inverted list attribute.
- OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 156-158
Chapter 3: SQL Application Tuning and Design
-

QUESTION 55:

In the CREATE TABLE syntax for an Index Organized Table, what is the purpose of the INCLUDING clause?

- A. It specifies the name of the primary key column in the index organized table.
- B. It specifies at which column to break a row into two pieces when a row's length exceeds the size set aside in TCTTHRESHOLD.
- C. It specifies what percentage of the entire data block to hold open in order to store the row data associated with a primary key value.
- D. It specifies the tablespace where the second half of the row data will be stored when the row's length exceeds the size set aside in PCTTHRESHOLD.

Answer: B

Explanation:

The INCLUDING clause specifies at which column to break a row into two pieces when a row's length exceeds the size set aside in PCTTHRESHOLD.

Incorrect Answers

A: It does not specify the name of the primary key column.

C: The PCTTHRESHOLD clause specifies what percentage of the entire data block to hold open in order to store the row data associated with a primary key value, which must be between 0 and 50 (default 50).

D: The OVERFLOW TABLESPACE clause specifies the tablespace where the second half of the row data will be stored when the row's length exceeds the size set aside in PCTTHRESHOLD.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 156-160
Chapter 3: SQL Application Tuning and Design

QUESTION 56:

What is the effect of setting the initialization parameter
QUERY_REWRITE_INTEGRITY to STALE_TOLERATED?

- A. Oracle server allows query rewrites based on declared, but not enforced, relationships.
- B. Oracle server allows all updated materialized views and constraints with RELY flag to be used for the query rewrites.
- C. Query rewrites can occur only if the server can guarantee consistency, and only up-to-date materialized views are used for the query rewrites.
- D. Query rewrites can occur even when the materialized view's data has not been refreshed and is inconsistent with the underlying detail data in the base tables.

Answer: D

Explanation:

If you will set the initialization parameter QUERY_REWRITE_INTEGRITY to STALE_TOLERATED, query rewrites can occur even when the view's data and the underlying table data are not current.

Incorrect Answers

A: This state of QUERY_REWRITE_INTEGRITY does not mean that Oracle server allows query rewrites based on declared, but not enforced, relationships. This can be done if TRUSTED state set for the QUERY_REWRITE_INTEGRITY parameter: query rewrites can occur when declared relationships exist, but without complete data currency.

B: Oracle server does not allow all updated materialized views and constraints with RELY flag to be used for the query rewrites if QUERY_REWRITE_INTEGRITY = STALE_TOLERATED.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 140-146
Chapter 3: SQL Application Tuning and Design

QUESTION 57:

What are two benefits of storing each table and index partition in a separate tablespace?
(Choose two)

- A. You can backup and recover each partition independently.

1Z0-033

- B. You can add and delete columns to partitions independently without affecting all the partitions.
- C. You can control the mapping of partitions to disk drives, which is important for balancing I/O LOAD.
- D. You can add and delete column constraints to partitions independently without affecting all the partitions.
- E. You can change a column data type in each partition independently without affecting all the other partitions.

Answer: A, C

Explanation:

Using partitions you can backup and recover each partition independently. Also you can control location of partitions on disk drives to improve I/O balancing.

Incorrect Answers

B: You cannot add or delete columns to partitions independently.

D: You cannot add or delete constraints to partitions independently.

E: You cannot change a column data type in each partition independently

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 162-181

Chapter 3: SQL Application Tuning and Design

QUESTION 58:

At the database instance level, how many database resource plans can be active at the same time?

- A. Only one
- B. All
- C. None
- D. Up to 2
- E. Up to 4
- F. Up to 8

Answer: A

Explanation:

Oracle8i introduced a new feature, Resource Manager, which was designed to improve the allocation and management of server resources needed by application users. In Oracle 9i, it is possible to control numerous aspects of application processing via Resource Manager. The database can have many resources plans, but only one can be active at any one time.

Incorrect Answers

B: The database can have many resources plans, but only one can be active at any one time.

C: Only one, not none, resource plan can be active at the same time.

D: Only one, not up to 2, resource plan can be active at the same time.

E: Only one, not up to 4, resource plan can be active at the same time.

F: Only one, not up to 8, resource plan can be active at the same time.

OCP: Oracle 9i Performance Tuning Study Guide, Jaseph C. Johnson, p. 517-547

Chapter 10: Operating System Tuning

QUESTION 59:

Which three statements are true regarding the effect of frequent checkpointing on your database? (Choose three)

- A. Can slow down a commit.
- B. Can affect instance recovery time.
- C. Can cause the CKPT process to hang.
- D. Can be influenced by redo log file size.
- E. Can case unexpected waits during redo log switches.

Answer: B, D, E

Explanation:

Frequent checkpoints will decrease instance recovery time. It can be influenced by redo log size: smaller it is - more frequent will be checkpoints. Checkpoint can be incomplete when log file is switched. When this occurs, the in-progress checkpoint is abandoned and a new checkpoint begun. Because incomplete checkpoints cause excess I/O that do not provide any recovery benefits, frequent occurrences of this event indicate that checkpoint activity should be tuned.

Incorrect Answers

A: Frequent checkpoints cannot slow down a commit.

C: They cannot cause the CKPT process to hang.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 338-340

Chapter 7: Tuning Redo Mechanisms

QUESTION 60:

You have a 512-byte OS block size. You want to cause a checkpoint event to occur every time 10K of data has been written from the Redo Log Buffer to the online redo log.

Which initialization parameter setting will achieve this?

- A. LOG_BUFFER=10240000
- B. LOG_CHECKPOINT_BYTES=10
- C. LOG_CHECKPOINT_TIMEOUT=10
- D. LOG_CHECKPOINT_INTERVAL=20
- E. LOG_CHECKPOINT_INTERVAL=10
- F. LOG_CHECKPOINT_TIMEOUT=5120

Answer: D

Explanation:

LOG_CHECKPOINT_INTERVAL initialization parameter set to 20 will achieve this goal: it shows number of OS blocks (20 * 0.5 K) when checkpoint will occur.

LOG_CHECKPOINT_INTERVAL specifies the frequency of checkpoints in terms of the number of redo log file blocks that can exist between an incremental checkpoint and the last block written to the redo log. This number refers to physical operating system blocks, not database blocks.

Incorrect Answers

A: LOG_BUFFER specifies the amount of memory, in bytes, that Oracle uses when buffering redo entries to a redo log file. Redo log entries contain a record of the changes that have been made to the database block buffers. The LGWR process writes redo log entries from the log buffer to a redo log file.

B: There is no LOG_CHECKPOINT_BYTES parameter in Oracle.

C: LOG_CHECKPOINT_TIMEOUT specifies the amount of time, in seconds, that has passed since the incremental checkpoint at the position where the last write to the redo log (sometimes called the tail of the log) occurred. This parameter also signifies that no buffer will remain dirty (in the cache) for more than integer seconds.

E: LOG_CHECKPOINT_INTERVAL value refers to physical operating system blocks, not database blocks.

F: LOG_CHECKPOINT_TIMEOUT specifies the amount of time, in seconds, that has passed since the incremental checkpoint at the position where the last write to the redo log (sometimes called the tail of the log) occurred. This parameter also signifies that no buffer will remain dirty (in the cache) for more than integer seconds.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 342-348

Chapter 7: Tuning Redo Mechanisms

QUESTION 61:

You just created a resource plan and placed this line in the init.ora

```
RESOURCE_MANAGER_PLAN = day_oltp
```

What does 'day_oltp' specify?

- A. Resource plan.
- B. Plan directive.
- C. Consumer group.
- D. Resource manager privilege.

Answer: A

Explanation:

RESOURCE_MANAGER_PLAN initialization parameter specifies a resource plan.

Incorrect Answers

B: It does not specify a plan directive. A plan directive is used to link a resource consumer group to a resource plan.

C: RESOURCE_MANAGER_PLAN initialization parameter does not specify a consumer group.

D: Resource manager privilege cannot be set with this initialization parameter.
OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 517-537
Chapter 10: Operating System Tuning

QUESTION 62:

Why do these steps eliminate row migration?

- Step 1: Run ANALYSE TABLE LIST CHAINED ROWS command
- Step 2: Copy the rows to another table
- Step 3: Delete the rows from the original table
- Step 4: Insert the rows from step 2 back into the original table

- A. Migration only occurs during an UPDATE operation.
- B. The migrated rows are removed with the DELETE command.
- C. Migration is automatically removed with the ANALYZE command.

Answer: A

Explanation:

These steps will eliminate row migration because it can happen only when inserted row is updated.

Incorrect Answers

B: The migrated rows will be removed with the DELETE command, but not only this step will allow you to avoid row migration.

C: The ANALYZE command does not remove a migration itself. It populates the CHAIN_CNT column of DBA_TABLES, which is otherwise null.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 404-405
Chapter 8: Tuning Disk I/O

QUESTION 63:

After running a query using V\$DISPATCHER, you increase the number of dispatchers. What would cause you to take this action?

- A. Users are waiting on a listener process.
- B. Users are waiting in dispatch processes.
- C. Users are waiting on shared server processes.
- D. Users are waiting on their dedicated connection process.

Answer: B

Explanation:

Like Shared Servers, the performance of the Dispatcher process can also be monitored using a "busy ratio". This ratio can be calculated using the V\$DISPATCHER dynamic performance view. Oracle recommends adding Dispatcher process if the Dispatcher busy rate consistently exceeds 50 percent. Another indication of Shared Server performance is

the amount of time that User Process are waiting to have their requests accepted by the Dispatchers.

Incorrect Answers

A: Listener process waiting have nothing to do with dispatchers.

C: If users are waiting on shared server processes you will not fix problem by increasing the number of dispatchers. Number of shared servers need to be increased instead.

D: Number of dispatchers is not related with dedicated connection process.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 300-302

Chapter 6: Tuning Other SGA Areas

QUESTION 64:

Which three statements regarding the SECONDS_IN_WAIT value for the log buffer space event in the V\$SESSION_WAIT view are true? (Choose three)

A. A SECONDS_IN_WAIT value close to zero is ideal.

B. A nonzero value in the SECONDS_IN_WAIT may indicate disk I/O contention on the redo log files.

C. The SECONDS_IN_WAIT value of the log buffer space event indicated time spent waiting for space in the redo log buffer.

D. A nonzero value in the SECONDS_IN_WAIT may be an indication the redo log buffers are too large and log switchers are not occurring fast enough.

Answer: A, B, C

Explanation:

V\$SESSION_WAIT performance view shows how long, and for which events, individual user sessions have waited. If column SECONDS_IN_WAIT value close to zero it is ideal situation: session does not wait for an event. A nonzero value in the SECONDS_IN_WAIT may be an indication the redo log buffers are too large to log switchers to occur fast enough. A nonzero value in the SECONDS_IN_WAIT may indicate disk I/O contention on the redo log files.

Incorrect Answers

D: The SECONDS_IN_WAIT value of the log buffer space event does not indicate time spent waiting for space in the redo log buffer.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 329-330

Chapter 7: Tuning Redo Mechanisms

QUESTION 65:

Which two parameters significantly impact the manual stripe size of the data files? (Choose two)

A. DB_BLOCK_SIZE

B. REDO_LOG_BUFFERS

C. DB_BLOCK_BUFFERS

- D. DB_BLOCK_MAX_DIRT_TARGET
- E. DB_FILE_MULTIBLOCK_READ_COUNT

Answer: A, E

Explanation:

The DB_BLOCK_MULTIBLOCK_READ_COUNT init.ora parameter determines the maximum number of database blocks that are read in one I/O operation by a user's Server Process whenever a full table scan read operation is performed. The primary block size is set at database creation and is specified in bytes by the init.ora parameter DB_BLOCK_SIZE. These two parameters can significantly impact the manual stripe size of the data files.

Incorrect Answers

B: There is no REDO_LOG_BUFFERS initialization parameter in Oracle. The size of the Redo Log Buffer is determined by the init.ora parameter LOG_BUFFER.

C: DB_BLOCK_BUFFERS specifies the number of database buffers in the buffer cache. It is one of several parameters that contribute to the total memory requirements of the SGA of an instance.

D: There is no DB_BLOCK_MAX_DIRT_TARGET parameter in Oracle.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 385-388
Chapter 8: Tuning Disk I/O

QUESTION 66:

You configured your database to use the Oracle Shared Server, but you did not configure the large pool. Where in memory will Oracle store user information?

- A. PGA
- B. JAVA POOL
- C. LARGE POOL
- D. SHARED POOL
- E. SHARED SQL AREA
- F. SHARED POOL RESERVED AREA

Answer: D

Explanation:

It will be stored in SHARED POOL area if it is not LARGE POOL is configured.

Incorrect Answers

A: PGA will not be used for this purpose.

B: SHARED POOL area will be used to store user information, not JAVA POOL.

C: LARGE POOL is not configured as question states.

E: There is no SHARED SQL AREA in Oracle9i.

F: SHARED POOL area will be used to store user information, not SHARED POOL RESERVED AREA.

QUESTION 67:

Which view shows the number of full table scans?

- A. V\$SYSSTAT
- B. V\$FILESTAT
- C. V\$SESSION
- D. V\$DATAFILE

Answer: A

Explanation:

V\$SYSSTAT performance view will be helpful to find out the number of full table scans. Occurrences of the statistics redo log space requests, DBWR buffers scanned, and DBWR LRU scans in the V\$SYSSTAT view are also useful for measuring the performance of DBW0.

Incorrect Answers

B: This view contains information about file read/write statistics.

C: This view contains information about current sessions.

D: This view contains datafile information from the control file.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 390-391
Chapter 8: Tuning Disk I/O

QUESTION 68:

Which two views would you query to monitor sessions related statistics? (Choose two)

- A. V\$SESSTAT
- B. V\$SESSION_EVENT
- C. V\$SESSION_STATS
- D. V\$SESSION_STATUS
- E. V\$WAITS_PER_SESSION

Answer: A, B

Explanation:

V\$SESSTAT view lists user session statistics. To find the name of the statistic associated with each statistic number (STATISTIC#), query the V\$STATNAME view.

V\$SESSION_EVENT view lists information on waits for an event by a session.

Incorrect Answers

C: There is no V\$SESSION_STATS view in Oracle.

D: There is no V\$SESSION_STATUS view in Oracle.

E: There is no V\$WAITS_PER_SESSION view in Oracle.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 386-388
Chapter 8: Tuning Disk I/O

QUESTION 69:

When a parallel query is used to perform a sort, what is the total amount of memory a factor to?

- A. SORT_AREA_SIZE * 2
- B. SORT_AREA_SIZE * degree of parallelism.
- C. SORT_AREA_SIZE *2* degree of parallelism.
- D. SORT_AREA_SIZE * divided up equally among the parallel query servers.
- E. SORT_AREA_SIZE * parallel query server take turns at using the memory.

Answer: C

Explanation:

The total amount of memory is SORT_AREA_SIZE *2* degree of parallelism.

Incorrect Answers

- A: This formula does not provide correct amount of memory. Parallelism can significantly improve a sorting time.
- B: SQL_AREA_SIZE should be multiplied on 2 to receive a correct formula.
- D: This formula incorrect: a total amount of memory for sorting is not related with query servers.
- E: This formula incorrect: a total amount of memory for sorting is not related with query servers.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 412-413
Chapter 8: Tuning Disk I/O

QUESTION 70:

How are deadlocks resolved within Oracle?

- A. The DBA must terminate the waiting session.
- B. The DBA must terminate the blocking session.
- C. Oracle detect deadlocks automatically and rolls back the statement which detects the deadlock.
- D. TopSessions monitors long running transactions and terminates any session which holds a lock longer than the limit specified by LOCK_THRESHOLD.

Answer: C

Explanation:

After detecting a deadlock Oracle rolls back the statement which detected the deadlock.

Incorrect Answers

- A: The DBA should not terminate the waiting session: Oracle can handle deadlock

situations automatically in most cases.

B: The DBA must terminate the blocking session: Oracle can handle deadlock situations automatically in most cases.

D: TopSessions does not terminate any session which holds a lock longer than the limit specified by LOCK_THRESHOLD.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 484-486
Chapter 9: Tuning Contention

QUESTION 71:

When a deadlock occurs, what should be used to diagnose the problem and determine how to prevent the deadlock from occurring again?

- A. Performance manager.
- B. ORA 00060 error message.
- C. The resulting trace file.

Answer: C

Explanation:

To diagnose and prevent a problem in the future the resulting trace file should be analyzed.

Incorrect Answers

A: Performance manager is used to tune database, but it will not be helpful to prevent deadlocks in the future: the resulting trace file can provide more valuable information.

B: This error only indicates that deadlock has been detected while waiting for resources, but it will not help you to prevent deadlocks or analyze their reason.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 484-486
Chapter 9: Tuning Contention

QUESTION 72:

What is a potential reason for a "snapshot too old" error message?

- A. You did not refresh your snapshots in time.
- B. An ITL entry in a data block has been reused.
- C. Are rollback segment extent sizes are too large.
- D. Your online redo log files are not big enough to snap your largest transactions.

Answer: B

Explanation:

If an ITL entry in a data block has been reused it can be a potential reason for a "snapshot too old" error message.

Incorrect Answers

A: This error does not related with Oracle snapshots itself. It deals with rollback

segment.

C: It can happen with small rollback segments.

D: This error is not related with online redo log files.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 445-446

Chapter 8: Tuning Disk I/O

QUESTION 73:

You are attempting to size the KEEP buffer pool and issue ANALYZE.... ESTIMATE STATISTIC command. Which three data dictionary tables should you query to obtain the total number of blocks required for an object? (Choose three)

- A. DBA_TABLES
- B. DBA_INDEXES
- C. DBA_SEGMENTS
- D. DBA_CLUSTERS

Answer: A, B, D

Explanation:

DBA_TABLES, DBA_INDEXES and DBA_CLUSTERS will provide you this information.

Incorrect Answers

C: DBA_SEGMENTS view works with segments, not tables or indexes.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 121-125

Chapter 3: SQL Application Tuning and Design

QUESTION 74:

The USERS table has thousands of rows and is accessed very often with an index (USERID_NDX) on a primary key (USER_ID).

Where should USERS and USERID_NDX be stored?

- A. Same tablespace.
- B. SYSTEM tablespace.
- C. Same tablespace on different disks.
- D. Different tablespace on different disks.

Answer: D

Explanation:

To avoid I/O contention you need to store data and indexes in different tablespaces on different disks.

Incorrect Answers

A: Storing data and indexes on the same tablespace can cause performance degradation.

B: Data and index data should not be stored on the same tablespace where SYSTEM

tablespace is stored.

C: Data and index data should be stored in different tablespaces.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 368-388

Chapter 8: Tuning Disk I/O

QUESTION 75:

You are creating a new rollback segment. Why should you choose the same value for the NEXT and INITIAL attributes?

- A. That depends on the PCTINCREASE value you specify.
- B. To avoid contention on the rollback segment header.
- C. Rollback segment extents are used in a circular way.
- D. Because you get an error message if you specify different values.

Answer: C

Explanation:

Like any other database segment, rollback segments are made up of extents, which are in turn comprised of five or more contiguous Oracle blocks. Within each rollback segment, Oracle uses the extents in a circular fashion until the rollback segment is full. That's why you should choose the same value for the NEXT and INITIAL attributes.

Incorrect Answers

A: You cannot set the PCTINCREASE value when creating a new rollback segment: it's always 0 for rollback segments.

B: Rollback segment extents are used in a circular way. Setting the same value for the NEXT and INITIAL attributes has nothing to do with contention on the rollback segment header.

D: You will get an error message, but this is not a reason, it's only a consequence of circular usage of rollback segments.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 425-429

Chapter 8: Tuning Disk I/O

QUESTION 76:

What should be your main redo log buffer tuning goal?

- A. Avoid space waits situations in the redo log buffer.
- B. Make the redo log buffer at least as large as the buffer cache.
- C. Mirror redo log files and store group members on different disks.
- D. Make sure that the LGWR process is faster than the database write processes.

Answer: A

Explanation:

You should avoid space waits situations in the redo log buffer.

Incorrect Answers

B: The redo log buffer should not be at least as large as the buffer cache.

C: Mirroring a redo log files and storing group members on different disks is not a tuning goal. You need to do this for better recoverability of the database.

D: If the LGWR process is faster than the database write processes redo logs will be populated with updates faster that DBWR will write dirty buffers on disk. This can cause checkpoint incomplete situations.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 324-327

Chapter 7: Tuning Redo Mechanisms

QUESTION 77:

You need some extra space in your database, so you issue the ALTER TABLE my_table DEALLOCATE UNUSED command. What happens?

A. All empty block of MY_TABLE are deallocated.

B. All blocks above the high-water mark of MY_TABLE are deallocated.

C. All blocks below the high-water mark of MY_TABLE are deallocated.

D. The high-water mark of MY_TABLE is recalculated and stored in the segment header.

Answer: B

Explanation:

Unused space can be released back to the tablespace using the ALTER TABLE tablename DEALLOCATE UNUSED command. The unused space above the High Water Mark (HWM) will be released.

Incorrect Answers

A: Not all empty blocks will be deallocated: only those above the high-water mark will be deallocated.

C: All blocks above, not below, the high-water mark of MY_TABLE are deallocated.

D: The high-water mark of MY_TABLE will not be recalculated after using this command.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 409-410

Chapter 8: Tuning Disk I/O

QUESTION 78:

The database is set up to run Multithreaded Server. Which view would show contention for server processes?

A. V\$QUEUE

B. V\$CIRCUIT

C. V\$SESSION

D. DBA_USERS

E. V\$CONNECTION

Answer: A

Explanation:

This view contains information on the multi-thread message queues. It shows information about contention for Shared Server Request and Response queues.

Incorrect Answers

B: V\$CIRCUIT view contains information about virtual circuits, which are user connections to the database through dispatchers and servers.

C: This view lists session information for each current session.

D: DBA_USERS data dictionary view provides information about database users.

E: There is no V\$CONNECTION data dictionary view in Oracle.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 298-300
Chapter 6: Tuning Other SGA Areas

QUESTION 79:

Which class of data describes to Oracle Expert how the database is used in daily operations?

- A. Schema class.
- B. Workload class.
- C. Instance class.
- D. Environment class.

Answer: B

Explanation:

Oracle Expert Workload class describes how the database is used in daily operations.

Incorrect Answers

A: Schema class does not describe how the database is used in daily operations.

C: Instance class does not describe how the database is used in daily operations.

D: Environment class does not describe how the database is used in daily operations.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 524-525
Chapter 10: Operating System Tuning

QUESTION 80:

On a database that is running Multithreaded server, which view would you query to get information for users with shared server connections?

- A. V\$CIRCUIT
- B. DBA_USERS
- C. DBA_CIRCUIT
- D. V\$DISPATCHER_USERS
- E. DBA_DISPATCHERS_USERS

Answer: A

Explanation:

V\$CIRCUIT view contains information about virtual circuits, which are user connections to the database through dispatchers and servers.

Incorrect Answers

B: DBA_USERS data dictionary view provides information about database users.

C: There is no DBA_CIRCUIT data dictionary view in Oracle.

D: There is no V\$DISPATCHER_USERS data dictionary view in Oracle.

E: There is no DBA_DISPATCHERS_USERS data dictionary view in Oracle.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 298-300

Chapter 6: Tuning Other SGA Areas

QUESTION 81:

To provide more free lists for a number of your database segments, what is one of your options?

A. Modify them with the INSERT_ _FREELIST command.

B. Drop and re-create them with the required FREELIST value.

C. Change the default storage parameter of the tablespace(s) where they are stored.

D. Modify the FREELIST_LIMIT parameter in your installation file and restart the instance.

Answer: B

Explanation:

To provide more free lists for a number of your database segments you can drop and re-create objects with the required FREELIST value. You should do this to avoid a Free List contention. If your application has many users performing frequent inserts, the application user's Server Process may experience waits when trying to access the Free List for a frequently inserted table.

Incorrect Answers

A: You cannot modify them with INSERT command.

C: You cannot to change them by altering the default storage parameter of the tablespace(s) where they are stored.

D: There is no FREELIST_LIMIT initialization parameter exists in Oracle.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 476-480

Chapter 9: Tuning Contention

QUESTION 82:

Which statement is true when connecting to the Oracle instance using the multithreaded server configuration?

A. The User Global Area (UGA) may only contain sort areas.

1Z0-033

- B. The User Global Area (UGA) may be accessible to dedicated servers.
- C. The User Global Area (UGA) components may reside in the large pool.
- D. The User Global Area (UGA) components may reside in the buffer cache.

Answer: C

Explanation:

The User Global Area (UGA) is only present in the Shared Pool if the Shared Server option is being used. In the shared server architecture, the User Global Area is used to cache application user session information. This information must be in shared location because the shared server architecture uses several different shared server processes to process SQL and PL/SQL activity against the instance. Since the shared server process that starts a transaction may not be the one that also finishes it, access to this session information must be stored in the UGA.

Incorrect Answers

- A: The User Global Area (UGA) may only contain not only sort areas.
 - B: In a non-shared (or dedicated) server situation, this session information is maintained in the application in the application user's private Process Global Area (PGA).
 - D: The User Global Area (UGA) components may not reside in the buffer cache.
- OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 209
Chapter 4: Tuning the Shared Pool
-

QUESTION 83:

Summing the system statistics, (db blocks gets + consistent gets) gives the total number of requests. What is the other system, statistic required to calculate the buffer cache hit ratio?

- A. Physical reads.
- B. Session logical gotten.
- C. Table scan blocks gotten.
- D. DBWR buffers scanned.

Answer: A

Explanation:

To calculate the buffer cache hit ratio you also need to know number of physical reads.

Incorrect Answers

- B: You don't need to know session logical gotten to calculate this ratio.
 - C: Table scan blocks gotten statistic is not required to calculate the buffer cache hit ratio.
 - D: DBWR buffers scanned also unnecessary to calculate it.
- OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 250-252
Chapter 5: Tuning the Database Buffer Cache
-

QUESTION 84:

Which statement about the LOG_BUFFER initialization parameter is true?

- A. The LOG_BUFFER parameter can be changed dynamically.
- B. The minimum value for the LOG_FILE parameter is 512KB.
- C. The LOG_BUFFER parameter value must be a multiple of the database block size.
- D. The LOG_BUFFER parameter value must be a multiple of the operating system block size.

Answer: D

Explanation:

The LOG_BUFFER parameter value should be a multiple of the operating system lock size. LOG_BUFFER specifies the amount of memory, in bytes, that Oracle uses when buffering redo entries to a redo log file. Redo log entries contain a record of the changes that have been made to the database block buffers. The LGWR process writes redo log entries from the log buffer to a redo log file.

Incorrect Answers

- A: The LOG_BUFFER parameter cannot be changed dynamically.
- B: The default, not minimum, value for the LOG_FILE parameter is 512KB. The minimum value is 64K.
- C: The LOG_BUFFER parameter value must be a multiple of the operating system, not database, block size.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 333-334
Chapter 7: Tuning Redo Mechanisms

QUESTION 85:

You determined that the values for REQUEST_FAILURES as seen from V\$SHARED_POOL_RESERVED is more than zero and always increasing. Which two actions would be appropriate? (Choose two)

- A. Decrease the value for LARGE_POOL_SIZE parameter.
- B. Increase the value for LARGE_POOL_SIZE parameter.
- C. Increase the value for SHARED_POOL_SIZE parameter.
- D. Decrease the value for SHARED_POOL_SIZE parameter.
- E. Increase the value for SHARED_POOL_RESERVED_SIZE parameter.
- F. Decrease the value for SHARED_POOL_RESERVED_SIZE parameter.

Answer: C, E

Explanation:

To avoid of increasing the value for REQUEST_FAILURES you should increase the value for SHARED_POOL_SIZE parameter and increase the value for SHARED_POOL_RESERVED_SIZE parameter. SHARED_POOL_RESERVED_SIZE specifies (in bytes) the shared pool space that is reserved for large contiguous requests for shared pool memory. You can use this parameter to avoid performance degradation in

the shared pool in situations where pool fragmentation forces Oracle to search for and free chunks of unused pool to satisfy the current request.

Incorrect Answers

A: You don't need to change the value for LARGE_POOL_SIZE parameter.

B: You don't need to change the value for LARGE_POOL_SIZE parameter.

D: You should increase, not decrease, the value for SHARED_POOL_SIZE parameter.

F: You should increase, not decrease, the value for SHARED_POOL_RESERVED_SIZE parameter.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 225-227

Chapter 4: Tuning the Shared Pool

QUESTION 86:

What are free lists used to identify?

A. Blocks available for inserts.

B. Free extents in a tablespace.

C. Blocks beyond the high water mark in a segment.

D. Segments belonging to a Parallel Server instance.

Answer: A

Explanation:

If your application has many users performing frequent inserts, the application user's Server Process may experience waits when trying to access the Free List for a frequently inserted table.

Incorrect Answers

B: Free Lists are not used to identify the free extents in a tablespace.

C: They do not represent blocks beyond the high water mark in a segment.

D: Segments belonging to a Parallel Server instance are not Free Lists.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 476-480

Chapter 9: Tuning Contention

QUESTION 87:

What should you confirm before changing the CURSOR_SPACE_FOR_TIME parameter in your initialization file to TRUE?

A. The TIMED_STATISTICS parameter is set to TRUE.

B. The hit percentage in the buffer cache is at least 95%.

C. The OPEN_CURSOR parameter is set to at least twice the default value.

D. The value in the RELOADS column of V\$LIBRARYCACHE is consistently zero.

Answer: D

Explanation:

1Z0-033

When `CURSOR_SPACE_FOR_TIME` init.ora parameter is set to `TRUE`, shared SQL areas are pinned in the Shared Pool. This prevents the LRU mechanism from removing a shared SQL area from memory unless all cursors that reference that shared SQL area are closed. This can improve the Library Cache hit ratio and reduce SQL execution time, at the expense of using more of the server's main memory. The default value for `CURSOR_SPACE_FOR_TIME` is `FALSE`. To switch it to `TRUE` you should confirm that the value in the `RELOADS` column of `V$LIBRARYCACHE` is consistently zero.

Incorrect Answers

A: You don't need to check that the `TIMED_STATISTICS` parameter is set to `TRUE`.

`TIMED_STATISTICS` specifies whether or not statistics related to time are collected.

B: It is not required to confirm that the hit percentage in the buffer cache is at least 95%.

C: The `OPEN_CURSOR` parameter setting to at least twice the default value is not necessary.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 232

Chapter 4: Tuning the Shared Pool

QUESTION 88:

Data dictionary information is held in memory longer than library cache data. Which is most likely to be true as a consequence of this?

A. You do not need to monitor library cache usage.

B. You have to tune the database buffer cache regularly.

C. You have to tune the library cache and dictionary cache independently.

D. Good hit ratios in the library cache imply acceptable hit ratios in the dictionary cache.

E. Good hit ratios in the dictionary cache imply acceptable hit ratios on the database buffer cache.

Answer: D

Explanation:

Good hit ratios in the library cache imply acceptable hit ratios in the dictionary cache because data dictionary information is held in memory longer than library cache data.

Incorrect Answers

A: You need to monitor the library and data dictionary cache usage.

B: You don't have to tune the database buffer cache regularly.

C: You should tune the library cache and dictionary cache consequently.

E: Good hit ratios in the dictionary cache do not prove acceptable hit ratios on the database buffer cache.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 218-220

Chapter 4: Tuning the Shared Pool

QUESTION 89:

User SCOTT creates an index with this statement:

```
CREATE INDEX emp_indx on employee(empno);
```

In which tablespace would the index be created?

- A. SYSTEM tablespace.
- B. Scott's default tablespace.
- C. Tablespace with rollback segments.
- D. Same tablespace as the EMPLOYEE table.

Answer: B

Explanation:

This index will be created in the Scott's default tablespace.

Incorrect Answers

A: If there is no default tablespace for the user Scott, index will be created in SYSTEM tablespace.

C: It will not be created in the ROLLBACK tablespace by default.

D: It will not be created in tablespace which contains the EMPLOYEE table, unless if it is Scott's default tablespace.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 422

Chapter 8: Tuning Disk I/O

QUESTION 90:

When setting multiple LRU latches in your initialization parameter file, what might you also consider setting?

- A. One buffer pool for each latch.
- B. One DBWn process for each latch.
- C. At one shared server for each latch.
- D. At least two DBWn processes for each latch.

Answer: B

Explanation:

You should set one DBWn process for each latch to avoid the latch contention.

Incorrect Answers

A: One DBWn process, not one buffer pool, should be set for each latch.

C: You don't need to consider to set shared servers for latches.

D: One, not two, DBWn process for each latch should be set.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 470-473

Chapter 9: Tuning Contention

QUESTION 91:

What can you use the values in the GETS and GETMISSES columns of V\$ROWCACHE to determine the hit ratio for?

1Z0-033

- A. Library cache.
- B. Dictionary cache.
- C. Entire shared pool.
- D. Large objects such as PL/SQL packages.

Answer: B

Explanation:

These values can be used to calculate the data dictionary cache hit ration. Formula $1 - \frac{\text{SUM}(\text{GETMISSES})}{\text{SUM}(\text{GETS})}$ will give you a value for the hit ratio.

Incorrect Answers

A: V\$ROWCACHE view is used only to calculate data dictionary cache hit ratio, not the library cache hit ratio.

C: Using V\$ROWCACHE view you cannot calculate the hit ratio for the entire shared pool.

D: You cannot determine the hit ratio for large objects such as PL/SQL packages using this values.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 218-219

Chapter 4: Tuning the Shared Pool

QUESTION 92:

You pinned an object in the shared pool using the DBMS_SHARED_POOL package. Which command could you use to unpin this object, assuming you are in a SQL *Plus session?

- A. ALERT SYSTEM FLUSH SHARED_POOL;
- B. EXECUTE dbms_shared_pool.unping;
- C. EXECUTE dbms_shared_pool.unpin;
- D. EXECUTE dbms_library_cache.unpin;

Answer: C

Explanation:

To unpin the object you should execute procedure UNKEEP of DBMS_SHARED_POOL package.

Incorrect Answers

A: This command will flash ENTIRE shared pool, not only unpin one object.

B: There is no UNPING procedure in DBMS_SHARED_POOL package.

D: There is no DBMS_LIBRARY_POOL package in Oracle.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 228-229

Chapter 4: Tuning the Shared Pool

QUESTION 93:

Which two parameters could result in problems when starting more shared servers?

(Choose two)

- A. PROCESSES
- B. MTS_MAX_SERVERS
- C. MTS_MAX_PROCESSES
- D. MTS_MAX_DISPATCHERS
- E. PARALLEL_MAX_SERVERS

Answer: A, B

Explanation:

PROCESSES and MTS_MAX_SERVERS init.ora parameters can be a limit when you try to start more shared servers.

Incorrect Answers

C: There is no MTS_MAX_PROCESSES init.ora parameter in Oracle.

D: MTS_MAX_DISPATCHERS specifies the maximum number of dispatcher processes allowed to be running simultaneously. The default value applies only if dispatchers have been configured for the system.

E: PARALLEL_MAX_SERVERS specifies the maximum number of parallel execution processes and parallel recovery processes for an instance. As demand increases, Oracle increases the number of processes from the number created at instance startup up to this value.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 307-308
Chapter 6: Tuning Other SGA Areas

QUESTION 94:

When does Oracle allocate memory for the large pool during instance startup?

- A. When the PARALELL_AUTOMATIC_TUNNING is set to FALSE.
- B. When Oracle is configured to use Multithreaded Server.
- C. When the LARGE_POOL_SIZE parameter is set to a valid value.
- D. When the large pool has a default value and is automatically allocated on instance startup.

Answer: C

Explanation:

When the LARGE_POOL_SIZE parameter is set to a valid value Oracle allocate memory for the large pool during instance startup. Once configured, the Oracle Server will automatically use the Large Pool to allocate memory to I/O server processes, RMAN, and Shared Server connections.

Incorrect Answers

A: By default, the LARGE_POOL_SIZE is zero (i.e., no Large Pool exist) unless the PARALLEL_AUTOMATIC_TUNING parameter is set to TRUE, in which case the size of the Large Pool is set automatically.

B: When Oracle is configured to use Multithreaded Server it does not mean that the Large Pool is set.

D: Oracle allocates memory for the large pool during instance startup, when the LARGE_POOL_SIZE parameter is set to a valid value, not when the large pool has a default value and is automatically allocated on instance startup.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 309
Chapter 6: Tuning Other SGA Areas

QUESTION 95:

Which procedure for the DBMS_RESOURCE_MANAGER package would first need to be performed when creating a new resource object?

- A. CREATE_PLAN
- B. CREATE_PENDING_AREA
- C. CREATE_CONSUMER_GROUP
- D. CREATE_PLAN_DIRECTIVE

Answer: B

Explanation:

Whenever a new resource consumer group, plan, or directive is created =, it is temporarily stored in the pending area until it is validated and written to the database. The purpose of the pending area is to give the DBA an opportunity to confirm that the definition of each consumer group, plan, and directive is correct before implementing it. The pending area is created using the DBMS_RESOURCE_MANAGER package and the CREATE_PENDING_AREA procedure.

Incorrect Answers

A: CREATE_PLAN procedure should be run after you created resource consumer group.

C: CREATE_CONSUMER_GROUP should be run after you created a pending area.

D: Directives can be created after you created resource plan.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 527-531

Chapter 10: Operating System Tuning

QUESTION 96:

What can database resource manager help you to limit for a set of users?

- A. Amount of I/O performed.
- B. Maximum connection time.
- C. Number of concurrent sessions.
- D. Number of Parallel Query servers available.

Answer: D

Explanation:

Resources that can be controlled through the use of Resource Manager include the amount of CPU allocated to each user who is a member of resource consumer group; the degree of parallelism allowed for Parallel Queries performed by a user who is a member of the resource consumer group; the amount of undo segment space that a user who is a member of a resource group is allowed to consume when performing a long running transaction; the total number of active, concurrent sessions that a given resource consumer group is allowed to have at one time ;the maximum expected total time of a database action taken by a user who is a member of the resource consumer group.

Incorrect Answers

A: It does not control an amount of I/O performed.

B: Database Resources Manager does not work with maximum connection time.

C: Number of concurrent sessions is not a resource controlled by Database Resources Manager.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 518-519

Chapter 10: Operating System Tuning

QUESTION 97:

Which single dynamic views is the most useful for determining buffer cache performance when using multiple buffer pools?

A. V\$SYSSTAT

B. V\$BUFFER_POOL

C. V\$SYSTEM_EVENT

D. V\$BUFFER_POOL_STATISTICS

Answer: D

Explanation:

V\$BUFFER_POOL_STATISTICS displays information about all buffer pools available for the instance. The "sets" pertain to the number of LRU latch sets.

Incorrect Answers

A: The V\$SYSSTAT dynamic performance view contains statistics regarding overall system performance, gathering since instance startup.

B: This view displays information about all buffer pools available for the instance. The "sets" pertain to the number of LRU latch sets.

C: This view contains information on total waits for an event. Note that the TIME_WAITED and AVERAGE_WAIT columns will contain a value of zero on those platforms that do not support a fast timing mechanism. If you are running on one of these platforms and you want this column to reflect true wait times, you must set TIMED_STATISTICS to TRUE in the parameter file ;doing this will have a small negative effect on system performance.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 275-277

Chapter 5: Tuning the Database Buffer Cache

QUESTION 98:

Which three types of tuning session scopes can Oracle Expert provide (Choose three)

- A. Session
- B. Instance
- C. Structure
- D. Application
- E. Operating system

Answer: B, C, D

Explanation:

Oracle Expert is a GUI tool that gathers and analyzes performance tuning data according to your specifications. Oracle Expert can provide instance, structure and application tuning session scopes.

Incorrect Answers

A: It does not tune session itself.

E: Oracle Expert cannot provide operating system tuning scope.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 72-74

Chapter 2: Sources of Tuning Information

QUESTION 99:

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

- A. V\$SYS_EVENTS
- B. V\$SYSTEM_EVENT
- C. V\$SESSION_WAIT
- D. V\$SYSTEM_STATUS

Answer: B

Explanation:

This view contains information on total waits for an event. Note that the TIME_WAITED and AVERAGE_WAIT columns will contain a value of zero on those platforms that do not support a fast timing mechanism. If you are running on one of these platforms and you want this column to reflect true wait times, you must set TIMED_STATISTICS to TRUE in the parameter file ;doing this will have a small negative effect on system performance.

Incorrect Answers

A: There is no data dictionary view V\$SYS_EVENTS in Oracle.

C: V\$SESSION_WAIT data dictionary view lists information about the CURRENT wait status of each session.

D: There is no data dictionary view V\$SYSTEM_STATUS in Oracle.

QUESTION 100:

Which statement is true when evaluating the buffer cache hit ratio?

- A. Minimizing physical reads will improve the buffer cache hit.
- B. The buffer cache hit ratio is unaffected by data or application design.
- C. The buffer cache hit ratio will improve the use of full table scans.
- D. The buffer cache hit ratio will always improve when the number of db block buffers in the SGA is increased.

Answer: A

Explanation:

When evaluating the buffer cache hit ratio you should try to minimize physical reads. Physical Reads statistic indicates the number of data blocks read from disk into the Buffer Cache since instance startup.

Incorrect Answers

- B: The buffer cache hit ratio is affected by data or application design.
- C: The buffer cache hit ratio will not improve the use of full table scans.
- D: The buffer cache hit ratio will always improve till some limit when the number of db block buffers in the SGA is increased.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 250-252
Chapter 5: Tuning the Database Buffer Cache

QUESTION 101:

Which three guidelines should you follow regarding the performance of redo logs?
(Choose three)

- A. Avoid using RAID 5 for redo logs.
- B. Place redo logs on disks without any other files.
- C. Manually stripe redo log files across several disks.
- D. Place redo log files on disks with non-Oracle files.
- E. Place members of the same group on different physical disks.

Answer: A, B, E

Explanation:

To improve the performance of redo logs you should avoid using RAID 5 for redo logs due to a decreasing performance after datafiles stripping. You need to try to place redo logs on disks without any other files. To increase recoverability of the database you need to place members of the same group on different physical disks.

Incorrect Answers

C: You should also avoid to stripe redo log files manual across several disks.

D: It's not required to place redo log files on disks with non-Oracle files. This will not improve performance. Better to place them on disks without any other files.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 324-327

Chapter 7: Tuning Redo Mechanisms

QUESTION 102:

What are two benefits of using locally managed tablespaces with segment space management set to auto? (Choose two)

- A. It eliminates the need to set PCTFREE.
- B. It eliminates any need to set PCTFREE or FREELISTS.
- C. It eliminates any need to set PCTUSED or FREELISTS.
- D. It improves the speed of space allocation and deallocation.
- E. It allows the DBA to use DB_FILE_MULTIBLOCK_READ_COUNT when setting extent sizes for a table.

Answer: C, D

Explanation:

By using locally managed tablespaces with segment space management set to auto you don't need to worry about PCTUSED or FREELISTS parameters. It improves also the speed of space deallocation and allocation.

Incorrect Answers

A: It does not eliminate the need to set PCTFREE parameter: you only don't need to set PCTUSED and FREELISTS parameters.

B: It eliminates any need to set PCTUSED or FREELISTS parameters, not PCTFREE parameter.

E: It does not allow the DBA to use DB_FILE_MULTIBLOCK_READ_COUNT.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 377

Chapter 8: Tuning Disk I/O

QUESTION 103:

Which two statements are true about checkpointing in Oracle? (Choose two)

- A. Small log files result in decreased checkpoint activity.
- B. Small log files result in increased checkpoint activity.
- C. Bigger temporary tablespaces favor increased checkpointing activity.
- D. If fast instance recovery is more important to you than achieving optimal run-time performance, then increase the checkpointing activity.
- E. If fast instance recovery is more important to you than achieving optimal run-time performance, then decrease the checkpointing activity.

Answer: B, D

Explanation:

Small log files result in increased checkpoint activity because they will be populated faster than large files.

To achieve the fast instance recovery time you should increase the checkpoint activity by increasing the FAST_START_MTTR_TARGET parameter, for example.

Incorrect Answers

A: Small log files result in increased checkpoint activity.

C: Temporary tablespaces do not interfere with checkpoint activity.

E: If fast instance recovery is more important to you than achieving optimal run-time performance, then increase the checkpointing activity.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 336-338

Chapter 7: Tuning Redo Mechanisms

QUESTION 104:

Assume that these are the current settings for your Oracle database:

1. The database has 3 Redo Log Groups with 1 member each of size 2 MB.
2. The LOG_CHECKPOINT_INTERVAL parameter is set to 10000.
3. The LOG_CHECKPOINT_TIMEOUT parameter is set to 1800.

Which action would increase the number of checkpoints in your database?

- A. Change the LOG_CHECKPOINT_TIMEOUT parameter to 2400 and restart the database.
- B. Change the LOG_CHECKPOINT_INTERVAL parameter to 12000 and restart the database.
- C. Increase the size of all the Redo Log Group member files to 3 MB and restart the database.
- D. Decrease the size of all the Redo Log Group member files to 1 MB and restart the database.

Answer: D

Explanation:

By decreasing the size of all the Redo Log Group member files to 1 MB and restarting the database you will increase the number of checkpoints in your database because the small Redo Log file requires more frequent checkpoints.

Incorrect Answers

A: Changing the LOG_CHECKPOINT_TIMEOUT parameter to 2400 will not increase the number of checkpoints in your database. LOG_CHECKPOINT_TIMEOUT specifies the amount of time, in seconds, that has passed since the incremental checkpoint at the position where the last write to the redo log (sometimes called the tail of the log) occurred. This parameter also signifies that no buffer will remain dirty (in the cache) for more than integer seconds.

B: LOG_CHECKPOINT_INTERVAL specifies the frequency of checkpoints in terms of the number of redo log file blocks that can exist between an incremental checkpoint and

the last block written to the redo log. This number refers to physical operating system blocks, not database blocks.

C: Increasing the size of all the Redo Log Group member files to 3 MB you will decrease a checkpoint activity.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 336-348
Chapter 7: Tuning Redo Mechanisms

QUESTION 105:

On a database that is running the Oracle Shared Server, which view would you query to get information for users with shared server connections?

- A. V\$CIRCUIT
- B. DBA_USERS
- C. DBA_CIRCUIT
- D. V\$DISPATCHER_USERS
- E. DBA_DISPACTHERS_USERS

Answer: A

Explanation:

This view contains information about virtual circuits, which are user connections to the database through dispatchers and servers.

Incorrect Answers

B: DBA_USERS data dictionary view provides information about database users, but not related with shared server connections.

C: This view does not exist in Oracle.

D: This view does not exist in Oracle.

E: This view does not exist in Oracle.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 298-300
Chapter 6: Tuning Other SGA Areas

QUESTION 106:

Which two statements are true regarding partitioned tables and indexes? (Choose two)

- A. A table which is part of a cluster can be partitioned.
- B. A global bitmap index can be created for a partitioned table.
- C. A table containing LONG or LONG RAW datatypes can be partitioned.
- D. A partitioned table can have partitioned and/or non-partitioned indexes.
- E. Privileges for partitions are granted on the parent table or index and not on individual partitions.
- F. Rule Based Optimization is used by default when a SQL statement accesses a partitioned table or index.

Answer: D, E

Explanation:

A partitioned table can have partitioned and/or non-partitioned indexes. Privileges for partitions are granted on the parent table or index and cannot be granted on individual partitions.

Incorrect Answers

A: It's a limitation that a table which is part of a cluster cannot be partitioned.

B: A global bitmap index can be created for a partitioned table.

C: A table containing LONG or LONG RAW datatypes cannot be partitioned.

F: Cost Based Optimization is used by default when a SQL statement accesses a partitioned table or index.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 165-181
Chapter 3: SQL Application Tuning and Design

QUESTION 107:

In a dictionary-managed tablespace, the SMON background process periodically coalesces neighboring free extents when which condition is true?

- A. When there is more than one data file in the tablespace.
- B. When the value of PCTINCREASE for the tablespace is zero.
- C. When the value of PCTINCREASE for the tablespace is NOT zero.
- D. When the value of INITIAL_EXTENT for the tablespace is greater than 1 MB.

Answer: C

Explanation:

The SMON background process periodically coalesces neighboring free extents when the value of PCTINCREASE for the tablespace is NOT zero.

Incorrect Answers

A: This condition is not required for the SMON background process to coalesce free extents.

B: When the value of PCTINCREASE for the tablespace is zero it will not coalesce free extents.

D: Size of the initial extent does not interfere with SMON process coalescing activity.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 400-401
Chapter 8: Tuning Disk I/O

QUESTION 108:

When a checkpoint occurs, the CKPT process _____.

- A. Writes the dirty buffers to the data files.
- B. Cleans up temporary segments that are no longer in use.
- C. Writes the redo entries from the log buffer to redo log files.
- D. Updates the control files to record the details of the checkpoint.

Answer: D

Explanation:

When a checkpoint occurs, the CKPT process updates the control files to record the details of the checkpoint. When an incremental checkpoint occurs, the CKPT background process updates the control file headers immediately. When a full checkpoint even occurs, database control files and database headers are updated by the CKPT background process to indicate that the checkpoint event has occurred.

Incorrect Answers

A: It does not write the dirty buffers to the data files. The DBWR process does this.

B: The CKPT process does not clean up temporary segments that are no longer in use.

C: LGWR process writes the redo entries from the log buffer to redo log files.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 336-338

Chapter 7: Tuning Redo Mechanisms

QUESTION 109:

You are responsible for selecting the Oracle optimizer mode for a new database that will be used in conjunction with a new Accounts Receivable application. The new application comes pre-configured with table and indexes statistics.

Which three Oracle optimizer modes might be suitable for the new application?

- A. RULE
- B. FIRST_ROWS_n
- C. ALL_ROWS
- D. FIRST_ROWS

Answer: B, C, D

Explanation:

When RULE is specified, statistics are not used. The CBO is a statistics-driven optimizer, and mode ALL_ROWS, FIRST_ROWS, FIRST_ROWS_n, CHOOSE are CBO settings.

Incorrect Answers

A: The rule-based optimizer mode does not use table or index statistics at all.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 99-101

Chapter 3: SQL Application Tuning and Design

QUESTION 110:

You need to determine what is the best multiplex redo log file configuration.

These facts describe your situation:

You have 2 independent disks: 'Disk-A' and 'Disk-B'.

You have 2 Redo Log Groups: 'Group-1' and 'Group-2'.

You have 2 Members in each Group.

Log1a and Log1b are both Members of Group-1.

Log2a and Log2b are both Members of Group-2.
Which is a recommended multiplexed redo log file configuration?

- A. Disk-A Disk-B
Log1a Log2a
Log1b Log2b
- B. Disk-A Disk-B
Log2a Log1a
Log1b
Log2b
- C. Disk-A Disk-B
Log1a Log 1b
Log2b Log2a
- D. Disk-A Disk-B
Log1a Log2b
Log2a
Log1b

Answer: C

Explanation:

This answer gives the best multiplexing redo log file configuration: each member of group is placed on different disk.

Incorrect Answers

A: All members of each group are placed on the same disk: this will decrease performance and recoverability of the database.

B: This configuration places all members of second group on the one disk.

D: This configuration places all members of first group on the one disk.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 345-348
Chapter 7: Tuning Redo Mechanisms

QUESTION 111:

You are planning a new Oracle9i database. Which four guidelines for tablespace planning should you follow? (Choose four)

- A. Use an UNDO tablespace.
- B. Separate applications data by tablespace.
- C. Place large tables in their own tablespace.
- D. Reserve the SYSTEM tablespace for data dictionary objects.
- E. Limit tablespaces to SYSTEM, TOOLS, UNDO, TEMPORARY, USERS and INDEXES.
- F. Place indexes on the same tablespace as the table for which they are created.

Answer: A, B, C, D

Explanation:

You should use a UNDO tablespace, separate applications data by tablespace, place large tables in their own tablespace and reserve the SYSTEM tablespace for data dictionary objects.

Incorrect Answers

E: You don't need to limit yourself to this list of tablespaces. For example, you may be need DATA tablespace.

F: You should never place indexes on the same tablespace as the table for which they are created.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 367-370
Chapter 8: Tuning Disk I/O

QUESTION 112:

Which statement about locally managed tablespaces in Oracle9i is true?

- A. Only the SYSTEM tablespace can be locally managed.
- B. Locally managed tablespaces are necessary for parallel queries.
- C. Free space information is contained within the locally managed tablespace.
- D. Free space and allocation of extents in locally managed tablespaces is tracked only in the data dictionary.

Answer: C

Explanation:

Free space information is contained within the locally managed tablespace. A Locally Managed Tablespace (LMT) uses a bitmap stored in the header of each of the tablespace's Datafiles instead of using Free Lists in the data dictionary to manage the allocation of space within the tablespace.

Incorrect Answers

A: Not only the SYSTEM tablespace can be locally managed.

B: Locally managed tablespaces are not necessary for parallel queries

D: Free space and allocation of extents in locally managed tablespaces is tracked only in locally managed tablespace.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 377
Chapter 8: Tuning Disk I/O

QUESTION 113:

You have set FAST_START_MTTR_TARGET so that instance recovery will take 5 minutes. Which view offers information regarding the systems ability to match the target?

- A. UNDO\$
- B. V\$DATAFILE
- C. V\$FILESTAT

- D. V\$TARGETRBA
- E. V\$INSTANCE_RECOVERY

Answer: E

Explanation:

This view monitors the mechanisms available to users to limit recovery I/O. Once the FAST_START_MTTR_TARGET parameter is set, you use the V\$INSTANCE_RECOVERY view to monitor its effectiveness.

Incorrect Answers

A: There is no UNDO\$ data dictionary view in Oracle.

B: V\$DATAFILE view does not offer this information.

C: This view contains information about file read/write statistics.

D: The view V\$TARGETRBA is obsolete and replaced by the view V\$INSTANCE_RECOVERY.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 343-344
Chapter 7: Tuning Redo Mechanisms

QUESTION 114:

The EMPLOYEE_MASTER table was created with the NOLOGGING attribute and has these columns:

FIRST_NAME VARCHAR2(60);

EMPLOYEE_ID NUMBER

Which two SQL statements can use NOLOGGING mode? (Choose two)

A. DELETE FROM EMPLOYEE_MASTER;

B. INSERT INTO EMPLOYEE_MASTER VALUES('ROGER',2000);

C. CREATE TABLE EMPLOYEE AS SELECT*FROM EMPLOYEE_MASTER;

D. CREATE INDEX EMPLOYEE_MASTER_PK ON

EMPLOYEE_MASTER(EMPLOYEE_ID);

E. UPDATE EMPLOYEE_MASTER

SET FIRST_NAME = 'JOHN'

WHERE EMPLOYEE_ID=1000;

Answer: C, D

Explanation:

You can use the NOLOGGING mode to create index, to rebuild index and to create table as select from the other table.

Incorrect Answers

A: It cannot be used for DML operations itself. Only table itself can be switched to this mode and then DELETE command can be performed for the table.

B: It cannot be used for DML operations itself. Only table itself can be switched to this mode and then INSERT command can be performed for the table.

E: It cannot be used for DML operations itself. Only table itself can be switched to this

mode and than UPDATE command can be performed for the table.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 335-336

Chapter 7: Tuning Redo Mechanisms

QUESTION 115:

A latch is defines as _____.

- A. A mechanism used by the Oracle optimizer to run user queries in parallel.
- B. A mechanism to prevent data files that make up the database from auto-extending.
- C. A low-level serialization mechanism to protect shared data structures in the system global area (SGA)
- D. A mechanism used by the background process PMON for cleaning up the caches and freeing up resources the user processes were using.

Answer: C

Explanation:

A latch is defines as a low-level serialization mechanism to protect shared data structures in the system global area (SGA). Latches are used to protect access to Oracle memory structures. A latch is a specialized type of lock that is used to serialize access to a particular memory structure or serialize the execution of kernel code.

Incorrect Answers

A: This is not a mechanism used by the Oracle optimizer to run user queries in parallel.

B: This is not a mechanism to prevent data files that make up the database from auto-extending.

D: Latches are not used by background process PMON.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 468-470

Chapter 9: Tuning Contention

QUESTION 116:

Which latch would be required when dirty blocks are written to the disk or when a server process is searching for blocks to write to?

- A. Shared pool latch
- B. Library cache latch.
- C. Cache buffers chains latch.
- D. Cache buffers LRU chain latch.

Answer: D

Explanation:

The Cache Buffers LRU Chain latch is used to manage the blocks on the LRU List in the Database Buffer Cache. This latch is used when Database Writer writes dirty buffers to disk and when a user's Server Process searches the LRU list for a free buffer during a

disk read.

Incorrect Answers

A: The Shared Pool latch is used to protect access to the Shared Pool's memory structures.

B: Like the Shared Pool latch, frequent waits for the Library Cache latch also indicates a poorly tuned Shared Pool.

C: The Cache Buffers Chains latch is accessed by user Server Processes when they are attempting to locate a data buffer that is cached in the Database Buffer Cache.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 470-472

Chapter 9: Tuning Contention

QUESTION 117:

You investigated latch contention by using the statistics in V\$LATCH view and the STATSPACK report, and you determined that there is contention for shared pool latch and library cache latch.

Which option would help reduce your shared pool and library cache latch contention?

A. Increase the db_block_buffers initialization parameter value and restart the database.

B. Identify similar SQL statements that could be shared in your application and convert them into sharable SQL with bind variables, thereby reducing unnecessary parsing.

C. Create additional indexes and rewrite the SQL in the application to avoid full table scans.

D. Set the PRE_PAGE_SGA initialization parameter to 'true' and restart the instance to allow Oracle to read the entire SGA into memory at instance startup.

Answer: B

Explanation:

To reduce your shared pool and library cache latch contention you need to identify similar SQL statements that could be shared in your application and convert them into sharable SQL with bind variables, thereby reducing unnecessary parsing.

Incorrect Answers

A: You will not reduce your shared pool and library cache latch contention by increasing the DB_BLOCK_BUFFERS initialization parameter value. DB_BLOCK_BUFFERS specifies the number of database buffers in the buffer cache. It is one of several parameters that contribute to the total memory requirements of the SGA of an instance.

C: This will not reduce your shared pool and library cache latch contention.

D: PRE_PAGE_SGA determines whether Oracle reads the entire SGA into memory at instance startup. Operating system page table entries are then pre-built for each page of the SG

A. This setting can increase the amount of time necessary for instance startup, but it is likely to decrease the amount of time necessary for Oracle to reach its full performance capacity after startup.

QUESTION 118:

Which two are valid modes for a process requesting latches? (Choose two)

- A. Scheduled
- B. Deferred
- C. Immediate
- D. Pre-Emptive
- E. Willing-to-Wait

Answer: C, E

Explanation:

There are two valid modes for a process requesting latches: immediate and willing-to-wait. If the latch is an immediate latch, the process requesting the latch continues to carry out other processing directives instead of waiting for the latch to become available. If the latch is a willing-to-wait latch, the process requesting the latch will wait for a short period of time and then request a latch again, perhaps waiting several more times, until it successfully attains the requested latch.

Incorrect Answers

- A: This is not valid mode for a process requesting latches.
- B: This is not valid mode for a process requesting latches.
- D: This is not valid mode for a process requesting latches.

QUESTION 119:

If a willing-to-wait latch request is satisfied on the first attempt, which statistic is incremented?

- A. GETS
- B. SLEEPS
- C. MISSES
- D. IMMEDIATE_GETS
- E. IMMEDIATE_MISSES
- F. IMMEDIATE_SLEEPS

Answer: A

Explanation:

GETS shows the number of times a willing-to-wait was acquired without waiting.

Incorrect Answers

- B: SLEEPS presents the number of times a process had to wait before obtaining a willing-to-wait latch.
- C: MISSES column shows the number of times a willing-to-wait latch was not acquired and a wait resulted.
- D: IMMEDIATE_GETS gives the number of times an immediate latch was acquired without waiting.
- E: IMMEDIATE_MISSES presents the number of times an immediate latch was not acquired and a retry resulted.
- F: There is no IMMEDIATE_SLEEPS column in the V\$LATCH view.
- OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 468-470
Chapter 9: Tuning Contention
-

QUESTION 120:

Which script creates views that will tell you who is holding or waiting for which tables?

- A. CATPROC.sql
- B. CATLOCK.sql
- C. CATPERF.sql
- D. CATBLOCK.sql
- E. CATBPROC.sql

Answer: D

Explanation:

A CATBLOCK.sql script creates the lock views that UTLOCKT.sql script needs, so you must run it before running UTLOCKT.sql.

Incorrect Answers

- A: CATPROC.sql script is required to load the initial PL/SQL environment.
- B: There is no CATLOCK.sql script in Oracle9i.
- C: There is no CATPERF.sql script in Oracle9i.
- E: There is no CATBPROC.sql script in Oracle9i.
- OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 468-470
Chapter 9: Tuning Contention
-

QUESTION 121:

Which statement about the amount of undo generated is true?

- A. The amount is the same for any DML operation.
- B. Deletes are inexpensive, because only the ROWID must be stored.
- C. Inserts are inexpensive, because only the ROWID must be stored.
- D. Updates are inexpensive, because only the new column value must be stored.

Answer: C

Explanation:

It's true that inserts are inexpensive: because only the ROWID must be stored.

Incorrect Answers

A: The amount is not the same for any DML operations. Inserts are inexpensive, because only the ROWID must be stored. Update is expensive, because old value needs to be stored also with new value.

B: Deletes are expensive, because old value needs to be stored, not only the ROWID.

D: Update is expensive, because old value needs to be stored also with new value.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 424-429

Chapter 8: Tuning Disk I/O

QUESTION 122:

Why can you NEVER achieve a value of zero in the GETISSES column if V\$ROWCACHE?

A. The database buffer cache can never be empty.

B. Recursive SQL has to be reparsed each time it is used.

C. Object definition must be loaded into the shared pool following instance startup.

D. An object cannot be pinned in the shared pool with the DBMS_SHARED_POOL package until it has been at least once.

Answer: C

Explanation:

You NEVER achieve a value of zero in the GETISSES column if V\$ROWCACHE because you need at least one time after instance startup load objects definitions from the data dictionary.

Incorrect Answers

A: This reason is not correct.

B: Recursive SQL has not to be re-parsed each time it is used.

D: The reason of this has nothing to do with DBMS_SHARED_POOL package.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 218

Chapter 3: SQL Application Tuning and Design

QUESTION 123:

Which action could result in less frequent checkpoints?

A. Increasing the number of redo log groups.

B. Increasing the value of DB_BLOCK_SIZE parameter.

C. Decreasing the value of the REDO_LOG_BUFFERS parameter.

D. Increasing the value of the FAST_START_IO_TARGET parameter.

Answer: D

Explanation:

FAST_START_IO_TARGET (available only with the Oracle Enterprise Edition) specifies the number of I/Os that should be needed during crash or instance recovery. When you set this parameter, DBWn writes dirty buffers out more aggressively, so that the number of blocks that must be processed during recovery stays below the value specified in the parameter. However, this parameter does not impose a hard limit on the number of recovery I/Os. Under transient workload situations, the number of I/Os needed during recovery may be greater than the value specified in this parameter. In such situations, DBWn will not slow down database activity. Smaller values for this parameter result in faster recovery times. This improvement in recovery performance is achieved at the expense of additional writing activity during normal processing. Setting this parameter to 0 disables fast-start checkpointing, which is the mechanism that limits the number of I/Os that need to be performed during recovery. All other writing activity is unaffected.

Incorrect Answers

A: Increasing the number of redo log groups could not result in less frequent checkpoints.

B: The value of DB_BLOCK_SIZE parameter has nothing to do with frequency of checkpoints.

C: There is no REDO_LOG_BUFFERS initialization parameter in Oracle. The size of the Redo Log Buffer is determined by the init.ora parameter LOG_BUFFER.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 343-345
Chapter 7: Tuning Redo Mechanisms

QUESTION 124:

Which two statements about database blocks are true? (Choose two)

- A. OLTP environment prefer a large block size.
- B. Small block size results in more block contention.
- C. Sequential access to large amounts of data favors a large block size.
- D. You can reduce the number of block visits by packaging rows as closely as possible into blocks.
- E. To change the database block size, you must shut down the instance and perform a START RESETLOGS after you make the change.

Answer: C, D

Explanation:

If you access large amount of data in sequential order it's better if it will be a large block size to reduce I/O operations. You can also reduce the number of block visits by packaging rows as closely as possible into blocks.

Incorrect Answers

A: OLTP environment prefer a small block size due to extensive DML operations.

B: Large block size results in more block contention because with larger size it's more chances that two transactions will try to work with the same block.

E: You cannot change the database block size this way: there is no START

RESETLOGS command in Oracle.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 187-188

Chapter 3: SQL Application Tuning and Design

QUESTION 125:

Sometimes the LGWR process must wait because DBW has not compressed checkpointing a file. How do you identify this situation?

- A. Check the V\$SESSION_WAIT view for the 'log buffer space' event.
- B. Check the alert.log file for the message "CHECKPOINT NOT COMPLETE".
- C. Check the 'redo buffer allocation retries' statistic in the V\$SYSSTAT view.
- D. Check the 'log file switch (checkpoint complete)' event in the V\$SYSTEM_EVENT view.

Answer: B

Explanation:

To find out this situation you should check the alert.log file for the message "CHECKPOINT NOT COMPLETE". This event shows how often the online redo log switched from one log to the next, before the checkpoint from the previous log switch had time to complete. When this occurs, the in-progress checkpoint is abandoned and a new checkpoint is begun. Because incomplete checkpoints cause excess I/O that do not provide any recovery benefits, frequent occurrences of this event indicate that checkpoint activity should be tuned.

Incorrect Answers

A: V\$SESSION_WAIT performance view shows how long, and for which events, individual user sessions have waited.

C: The V\$SYSSTAT dynamic performance view contains statistics regarding overall system performance, gathering since instance startup.

D: This view contains information on total waits for an event.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 338-342

Chapter 7: Tuning Redo Mechanisms

QUESTION 126:

You want to reduce the amount of redo generated for your database. What are three ways to achieve this goal? (Choose three)

- A. Use NOLOGGING mode in SQL statements.
- B. Use direct load UPDATED to NOLOGGING mode.
- C. Use direct path loading without archiving.
- D. Use direct path loading with archiving using NOLOGGING mode.
- E. Start you instance with the NOLOGGING initialization parameter.

Answer: A, C, D

Explanation:

To reduce the amount of redo generated for your database you can use NOLOGGING mode in some SQL statements, use direct path loading without archiving or use direct path loading with archiving using NOLOGGING mode.

Incorrect Answers

B: There is no direct load UPDATE command with NOLOGGING mode in Oracle.

E: There is no NOLOGGING initialization parameter in init.ora file.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 335-336
Chapter 7: Tuning Redo Mechanisms

QUESTION 127:

The alert log file for a database instance indicated that the checkpoints are frequently failing to complete. Which action would be a remedy in this situation?

- A. Increase the number of archiver (ARCn) processes.
- B. Increasing the number of members for all log groups.
- C. Increasing the number of log writer (LGWR) processes.
- D. Increasing the number of database writer (DBWn) processes.

Answer: D

Explanation:

Sometimes the LGWR process must wait because DBW has not finished to checkpoint a file. To speed up data modifications writing to datafiles you should increase the number of database writer (DBWn) processes.

Incorrect Answers

A: This event has nothing to do with ARCn archiver processes.

B: You don't need to increasing the number of members for all log groups.

C: Increasing the number of log writer (LGWR) processes will not fix the problem: it may be make it worse because there will be more log writer LGWR processes that DBWn database writer processes.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 338-342
Chapter 7: Tuning Redo Mechanisms

QUESTION 128:

When the archive process encounters an error, which parameter determines the directory where trace files are written?

- A. UTL_FILE_DIR
- B. CORE_DUMP_DEST
- C. LOG_ARCHIVE_DEST
- D. BACKGROUND_DUMP_DEST

Answer: D

Explanation:

BACKGROUND_DUMP_DEST directory contains the trace file when the archive process encounters an error.

Incorrect Answers

A: UTL_FILE_DIR lets you specify one or more directories that Oracle should use for PL/SQL file I/O. If you are specifying multiple directories, you must repeat the UTL_FILE_DIR parameter for each directory on separate lines of the initialization parameter file.

B: CORE_DUMP_DEST is primarily a UNIX parameter and may not be supported on your platform. It specifies the directory where Oracle dumps core files.

C: LOG_ARCHIVE_DEST is applicable only if you are running the database in ARCHIVELOG mode or are recovering a database from archived redo logs.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 40-41
Chapter 2: Sources of Tuning Information

QUESTION 129:

Which two statements are true when you are using automatic undo management?
(Choose two)

A. If you do not specify the init.ora parameter UNDO_TABLESPACE when you start the database, Oracle will make use of all undo tablespaces within the database to store automatic undo segments.

B. If you do not specify the init.ora parameter UNDO_TABLESPACE when you start the database, Oracle will make use of the first undo tablespace that was created within the database to store automatic undo segments.

C. Oracle will extend the datafiles for the undo tablespace to meet the requirements of the time specified by the init.ora parameter UNDO_RETENTION.

Oracle will always maintain committed transactions for this amount of time.

D. Oracle will not extend the datafiles for the undo tablespace to meet the requirements of the time specified by the init.ora parameter UNDO_RETENTION.

Oracle will attempt to maintain committed transactions for this amount of time, but may not always be successful.

Answer: B, D

Explanation:

Answer B is true because Oracle will make use of the first undo tablespace that was created within the database to store automatic undo segments if you do not specify the init.ora parameter UNDO_TABLESPACE. Answer D is true, maintaining committed transaction is not always successful. Because maybe there is not enough free undo space for current active transaction and the system will reuse undo space which can lead ORA-01555 : Snapshot too old error.

Incorrect Answers

A: If you do not specify the init.ora parameter UNDO_TABLESPACE when you start the database, Oracle will make use of the first, not all, undo tablespace within the database to store automatic undo segments.

C: Oracle will NOT always maintain committed transactions for UNDO_RETENTION amount of time. In case of shortage of free undo space it can reuse undo space.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 338-342
Chapter 7: Tuning Redo Mechanisms

QUESTION 130:

Which type of change to an application is most likely to improve performance of the library cache?

- A. Adding more frequent COMMIT statements.
- B. Replacing bind variables with constraints.
- C. Reusing as much generic code as possible.
- D. Replacing database constraints with triggers.

Answer: C

Explanation:

If you are reusing as much generic code as possible it can improve significantly performance of the library cache.

Incorrect Answers

A: Adding more frequent COMMIT statements will not improve performance of the library cache.

B: You should use more bind variables than constraints to improve it.

D: Database constraints and triggers have nothing to do with the library cache.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 205-208
Chapter 4: Tuning the Shared Pool

QUESTION 131:

The cost-based optimizer can choose between a nested loops join and a sort merge join operation. All tables are analyzed and the OPTIMIZER_MODE is set to FIRST_ROWS. Which execution plan will be the result?

- A. The sort-merge join
- B. The nested loops join
- C. This depends on some sort parameter values.
- D. This depends on the number of rows in each table.

Answer: B

Explanation:

If the OPTIMIZER_MODE is set to FIRST_ROWS and all tables are analyzed the

cost-based optimizer will choose a nested loops join operation.

Incorrect Answers

A: If the OPTIMIZER_MODE is not set to FIRST_ROWS or tables are not analyzed the cost-based optimizer can choose a sort merge join operation.

C: It does not depend on some sort parameter values.

D: It does not depend on the number of rows in each table.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 99-101

Chapter 3: SQL Application Tuning and Design

QUESTION 132:

To control fragmentation of your shared pool space, when is the best time for you to pin objects with the DBMS_SHARED_POOL package?

A. Immediately after instance startup.

B. After an object has been used for the first time.

C. When the V\$LIBRARYCACHE view contains higher values in the RELOADS column than in the PINS column.

D. When the sum of values in the SHARABLE_MEMORY column of the V\$DB_OBJECT_CACHE view exceeds the value of the SHARED_POOL_SIZE initialization parameter.

Answer: A

Explanation:

The best time for you to pin objects with the DBMS_SHARED_POOL package immediately after instance startup.

Incorrect Answers

B: It will not be the best time after an object has been used for the first time.

C: It has nothing to do with V\$LIBRARYCACHE view values.

D: It has nothing to do with V\$DB_OBJECT_CACHE view values.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 228-229

Chapter 4: Tuning the Shared Pool

QUESTION 133:

You have been hired as a senior DBA at a site that has just started to use Oracle9i.

You need to develop a method to tune the database for best performance for all users. You must present a list of tuning goals that your job performance can be measured by.

What are three acceptable tuning goals? (Choose three)

A. Increase the database buffer hit ratio from 75% to 85%.

B. Reduce the amount of redo logging to a minimum during data loads.

C. Reduce the average instance recovery time from five minutes to two minutes.

D. Reduce the average response time for the order entry application from 15 seconds to

1Z0-033

as quickly as possible.

E. Increase the ratio GETHITS to GETS in the V\$LIBRARYCACHE view from the current 70 percent to 90 percent.

Answer: A, C, E

Explanation:

These three tuning goals present a list using that your job performance can be measured by.

Incorrect Answers

B: It's tuning goal also, but it does not provide you with target redo logging you should achieve.

D: This tuning goal also does not show target response time for the order entry application, so it cannot be used as acceptable tuning goal.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 6-10

Chapter 1: Introduction to Performance Tuning

QUESTION 134:

You should query the buffer cache hit ratio during normal working loads, but not immediately after startup. Why?

A. There are no statistics to query because there is no database activity.

B. Statistics are only available once the database has been active for a period of time.

C. Querying the buffer cache statistics immediately after database startup will yield incomplete results.

D. The buffer cache is empty when the instance starts, resulting in more physical reads after startup.

Answer: B

Explanation:

Statistics are only available once the database has been active for a period of time.

Incorrect Answers

A: There are some statistics to query, but it will be not very accurate due to recent database startup.

C: This statement is not correct: querying the buffer cache statistics immediately after database startup will present incomplete results.

D: The buffer cache even when the instance starts cannot be empty because at least some information from the dictionary tables should be loaded for normal database functioning.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 275-277

Chapter 5: Tuning the Database Buffer Cache

QUESTION 135:

Every afternoon processes that are running against the database slow down almost to a

1Z0-033

stop. A tool is needed to pinpoint which user processes are using the most database resources. Which Oracle Enterprise Manager tool will allow you to pinpoint which session is using the highest amount of database resources?

- A. TOPSESSIONS
- B. SQL ANALYZE
- C. ORACLE EXPERT
- D. PERFORMANCE MANAGER

Answer: A

Explanation:

All these tools can provide you with information you need for tuning purposes, but TopSessions Oracle Enterprise Manager tool will provide you with necessary information to pinpoint which session is using the highest amount of database resources.

Incorrect Answers

B: SQL Analyze tool can help you to tune SQL statements when you will find what session uses them and slow down the database performance.

C: Oracle Expert is tool for overall database tuning, not session level tool.

D: Performance Manager is also used to monitor and tune entire database disregarding the session information.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 40-41

Chapter 2: Sources of Tuning Information

QUESTION 136:

Which two statements are true regarding automatic undo segments? (Choose two)

- A. Undo tablespaces must be created as **DICTIONARY MANAGED**.
- B. Undo tablespaces must be created as **LOCALLY MANAGED**.
- C. Any tablespace can be used for automatic undo segment management. You specify which tablespace to use with the init.ora parameter **UNDO_TABLESPACE**.
- D. An undo tablespace can only be dropped if it is **NOT** the active undo tablespace for the database, and all transactions using that undo tablespace have committed or rolled back.

Answer: B, D

Explanation:

It's requirement that automatic undo segments must be created as **LOCALLY MANAGED**. It can be dropped if it is **NOT** the active undo tablespace for the database, and all transactions using that undo tablespace have committed or rolled back.

Incorrect Answers

A: Undo tablespaces must be created as **LOCALLY MANAGED**, not **DICTIONARY MANAGED**.

C: Not any tablespace can be used for automatic undo segment management.

QUESTION 137:

You discover that you have free buffer waits occurring in your Oracle9i database. You determine that DBWn performance is the problem. Your investigation shows that the I/O is well within the capabilities of your disks and disk controllers. The disks are not raw devices, there are multiple CPUs and you have Asynchronous I/O turned on.

What is the best solution?

- A. Configure DBWR_IO_SLAVES.
- B. Configure DB_WRITER_PROCESSES.
- C. Turn asynchronous I/O off and configure DBWR_IO_SLAVES.
- D. For best performance, configure both DBWR_IO_SLAVES and DB_WRITER_PROCESSES.

Answer: B

Explanation:

The DB_WRITER_PROCESSES initialization parameter lets you configure multiple database writer processes (from DBW0 to DBW9 and from DBWa to DBWj). Configuring multiple DBWR processes distributes the work required to identify buffers to be written, and it also distributes the I/O load over these processes. Multiple db writer processes are highly recommended for systems with multiple CPUs (at least one db writer for every 8 CPUs) or multiple processor groups (at least as many db writers as processor groups).

Incorrect Answers

A: If it is not practical to use multiple DBWR processes, then Oracle provides a facility whereby the I/O load can be distributed over multiple slave processes. The DBWR process is the only process that scans the buffer cache LRU list for blocks to be written out. However, the I/O for those blocks is performed by the I/O slaves. The number of I/O slaves is determined by the parameter DBWR_IO_SLAVES.

C: Turning asynchronous I/O off will decrease DBWn performance.

D: It is no sense to configure both DBWR_IO_SLAVES and DB_WRITER_PROCESSES simultaneously. DBWR_IO_SLAVES is intended for scenarios where you cannot use multiple DB_WRITER_PROCESSES (for example, where you have a single CPU). I/O slaves are also useful when asynchronous I/O is not available, because the multiple I/O slaves simulate nonblocking, asynchronous requests by freeing DBWR to continue identifying blocks in the cache to be written.

Asynchronous I/O at the operating system level, if you have it, is generally preferred.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 275-277
Chapter 5: Tuning the Database Buffer Cache

QUESTION 138:

What are three advantages of having a large database block size? (Choose three)

- A. Large database block sizes reduce block contention.
- B. Large database block sizes are good for random access.
- C. Large database block sizes are good for sequential access.
- D. Large database block sizes increase performance of index reads.
- E. Large database block sizes have less overhead, and more room is left to store data.

Answer: C, D, E

Explanation:

Large database block sizes are good for sequential access. They are increasing performance of index reads. And large database block sizes have less overhead, and more room is left to store data.

Incorrect Answers

A: They do not reduce block contention: small blocks are more useful to avoid block contention.

B: Large database block sizes are not good candidates for random access.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 445-446
Chapter 8: Tuning Disk I/O

QUESTION 139:

Which statement is true when connections to the Oracle instance use the shared server configuration?

- A. The User Global Area (UGA) does not include the sort areas.
- B. The User Global Area (UGA) resides in the private user memory.
- C. The User Global Area (UGA) area contains the private SQL areas.
- D. The User Global Area (UGA) is not accessible to all the shared servers.

Answer: C

Explanation:

The User Global Area (UGA) area contains the private SQL areas.

Incorrect Answers

A: It includes the sort areas.

B: The User Global Area (UGA) does not reside in the private user memory.

D: The User Global Area (UGA) is accessible to all the shared servers.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 300-302
Chapter 6: Tuning Other SGA Areas

QUESTION 140:

1Z0-033

Your database is used by an OLTP application.

1. What is the recommended Library Cache GETHITRATIO percentage you should try to achieve?
2. How can you increase the GETHITRATIO percentage in the V\$LIBRARYCACHE?

- A. 80%-85;analyze the table
- B. 85%-90;improve performance by increasing the SHARED_POOL_SIZE
- C. 90%-95;improve performance by increasing the LARGE_POOL_SIZE
- D. 90%-100;improve the efficiency of the application code

Answer: D

Explanation:

For OLTP application you should try to achieve 90% - 100% GETHITRATION percentage. You need to improve the efficiency of the application code.

Incorrect Answers

A: It's no sense to analyze the tables in an OLTP application: statistics will be obsolete very soon due to high volume of DML operations.

B: These figures does not provide good enough goal to achieve: GETHITRATION should not be less than 90%. Otherwise database performance will degrade.

C: You can improve performance by increasing the LARGE_POOL_SIZE. But this will be more valuable for Warehouse, not OLTP systems. Unlike the shared pool, the large pool does not have an LRU list. Oracle does not attempt to age objects out of the large pool. You should consider configuring a large pool if your instance uses any of the following: parallel query (parallel query uses shared pool memory to cache parallel execution message buffers), Recovery Manager (Recovery Manager uses the shared pool to cache I/O buffers during backup and restore operations. For I/O server processes and backup and restore operations, Oracle allocates buffers that are a few hundred kilobytes in size).

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 205-208
Chapter 4: Tuning the Shared Pool

QUESTION 141:

Which three actions avoid the disadvantages of dynamic extent allocation? (Choose three)

- A. Increase the PCTINCREASE on the table.
- B. Create locally managed tablespaces for the objects that extend continuously.
- C. Coalesce the tablespace to force neighboring free extents to be merged together.
- D. Monitor segments ready to extend and manually add the extents during off peak hours.
- E. Export and import all the tables in a tablespace to reorganize the blocks in the tables.
- F. Choose storage parameters that allocate extents large enough to accommodate all of your data when you create the object.

Answer: B, D, F

Explanation:

If you create locally managed tablespaces you still cannot control extent allocation process. By manually adding extents you can avoid dynamic extent allocation. If you choose storage parameters that allocate extents large enough to accommodate all of your data when you create the object it allows you to control the dynamic extent allocation.

Incorrect Answers

A: By increasing the PCTINCREASE parameter you just can be sure that next extent size larger will be larger than previous one.

C: Coalescing the tablespace you just merge free extents, not used.

D: By manually adding extents you still increasing amount of extents for table.

E: By exporting and importing all the tables in a tablespace you can compress amount of extents for each table to one, but it will not allow you to avoid the disadvantages of dynamic extent allocation.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 445-446
Chapter 8: Tuning Disk I/O

QUESTION 142:

What causes session waits in the log buffer?

A. The rollback tablespace extends.

B. Data has not been committed to the database.

C. A transaction has an exclusive lock on a row.

D. Two users update the same row at the same time.

E. Data is written into the log buffer faster than LGWR can write it out.

Answer: E

Explanation:

This issue is called redo log contention: data is written into the log buffer faster than LGWR can write it out.

Incorrect Answers

A: It has nothing to do with rollback tablespace extension.

B: It is not related with frequency of committing or data roll back.

C: Redo log contention cannot be affected by exclusive lock on a row.

D: This statement is incorrect.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 338-342
Chapter 7: Tuning Redo Mechanisms

QUESTION 143:

Which statement about the PCTFREE parameter is true?

A. If you expect zero update activity on your data, you should set the PCTFREE value to

100 to maximize data block utilization.

B. The PCTFREE parameter sets the minimum percentage of a data block that can be used for future inserts of new rows in that block.

C. A smaller PCTFREE is suitable for segments that are frequently updated, and will improve update performance.

D. The PCTFREE parameter sets the minimum percentage of a data block to be reserved for future updates to rows that already exist in that block.

Answer: D

Explanation:

The PCTFREE parameter sets the minimum percentage of a data block to be reserved for future updates to rows that already exist in that block.

Incorrect Answers

A: PCTFREE cannot be set to 100: it means that entire block should be reserved for updates. But it's impossible because PCFREE + PCTUSED < 100.

B: The PCTFREE parameter sets the minimum percentage of a data block to be reserved for future updates to rows that already exist in that block, not for future inserts of new rows in that block.

C: A larger PCTFREE is suitable for segments that are frequently updated, and will improve update performance: a smaller PCTFREE will decrease update performance.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 445-446

Chapter 8: Tuning Disk I/O

QUESTION 144:

The message "ORA - 29554: unhandled Java out of memory condition" is caused by which condition?

A. JAVA_MAX_SESSIONSPACE_SIZE is reached.

B. JAVA_SOFT_SESSIONSPACE_LIMIT is reached.

C. A user's session-duration state exceeds 1 M.

D. Java Pool 'memory in use' is equal to the JAVA_POOL_SIZE setting.

Answer: A

Explanation:

The session encountered an out of memory condition in Java from which it could not recover. Java session state was cleared. JAVA_MAX_SESSIONSPACE_SIZE is Java session space is the memory that holds Java state from one database call to another.

JAVA_MAX_SESSIONSPACE_SIZE specifies (in bytes) the maximum amount of session space made available to a Java program executing in the server. When a user's session-duration Java state attempts to exceed this amount, the Java virtual machine kills the session with an out-of-memory failure.

Incorrect Answers

B: JAVA_SOFT_SESSIONSPACE_LIMIT parameter shows a Java session space is the

memory that holds Java state from one database call to another.

JAVA_SOFT_SESSIONSPACE_LIMIT specifies (in bytes) a soft limit on Java memory usage in a session, as a means to warn you if a user's session-duration Java state is using too much memory. When a user's session-duration Java state exceeds this size, Oracle generates a warning that goes into the trace files.

C: A user's session-duration state exceeded 1 M cannot cause this error.

D: This error is not related to Java Pool 'memory in use' equality to the JAVA_POOL_SIZE setting.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 99-101
Chapter 3: SQL Application Tuning and Design

QUESTION 145:

You have been taking snapshots once a week over the last two years. You will use these snapshots to outline the cost associated with the increase in database resources over the past two years. Which tool or utility will allow you to compare the snapshots?

- A. STATSPACK
- B. LOCK MANAGER
- C. ORACLE EXPERT
- D. UTLBSTAT /UTLESTAT

Answer: A

Explanation:

The term snapshot as used here denotes a set of performance statistics gathered at a single time, identified by a unique ID that includes the snapshot number (or SNAP_ID). This kind of snapshot has nothing to do with Oracle's snapshot replication technology. Statspack is a set of SQL, PL/SQL, and SQL*Plus scripts that allow the collection, automation, storage, and viewing of performance data. Statspack differs from the existing UTLBSTAT/UTLESTAT performance scripts in the following ways:

1. Statspack collects more data, including high-resource SQL.
2. Statspack precalculates many ratios useful when performance tuning, such as cache hit ratios, rates, and transaction statistics. (Many of these ratios must be calculated manually when using BSTAT/ESTAT).
3. Permanent tables owned by PERFSTAT store performance statistics. Instead of creating/dropping tables each time, data is inserted into the preexisting tables. This makes it easier to compare historical data.
4. Statspack separates data collection from report generation. Data is collected when a snapshot is taken. The performance engineer then runs the performance report and views the data collected.
5. Data collection is easy to automate using either DBMS_JOB or an operating system utility.
6. Statspack considers a transaction to finish either with a COMMIT or a ROLLBACK, and so calculates the number of transactions as 'user commits' + 'user rollbacks.'

BSTAT/ESTAT considers a transaction to complete with a COMMIT only, and so

assumes that transactions = 'user commits.' For this reason, comparing statistics for each transaction between Statspack and BSTAT/ESTAT can result in significantly different ratios.

Incorrect Answers

B: Lock Manager is used to monitor locks only.

C: Oracle Expert will not provide you with functionality described in the question.

D: UTLBSTAT /UTLESTAT performance scripts present you with less valuable statistics.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 40-41

Chapter 2: Sources of Tuning Information

QUESTION 146:

Which two statements are true regarding tuning the Data Dictionary Cache? (Choose two)

A. The ratio of GETMISSES to GETS found in the V\$ROWCACHE view should be less than 5 percent.

B. If the hit ratio for the library cache is acceptable, the hit ratio for the data dictionary should also be acceptable.

C. On instance startup, the data dictionary cache contains no data, so any SQL statement is likely to result in cache misses.

D. On instance startup the data dictionary cache is automatically loaded with the most commonly used information, so many SQL statements will not cause cache misses.

Answer: B, C

Explanation:

Configuring the library cache to an acceptable size usually ensures that the data dictionary cache is also properly sized. So tuning Library Cache will tune Row Cache indirectly. When instance starts, there are no data in the data dictionary cache, so any SQL statement is likely to result in cache misses.

Incorrect Answers

A: The ratio of GETMISSES to GETS found in the V\$ROWCACHE view should be less than 1 percent.

D: On instance startup the data dictionary cache is not automatically loaded with the most commonly used information.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 205-208

Chapter 4: Tuning the Shared Pool

QUESTION 147:

When a deadlock situation is detected by Oracle, which three results occur? (Choose three)

A. A trace file is generated.

- B. An error is reported in the alert log.
- C. All sessions involved in the deadlock are terminated.
- D. All transactions currently in progress are terminated.
- E. All database sessions hang indefinitely because of the deadlock.
- F. The statement causing the deadlock in the session where the deadlock is detected is rolled back.

Answer: A, B, F

Explanation:

When a deadlock situation is detected by Oracle, a trace file is generated, an error is reported in the alert log and the statement causing the deadlock in the session where the deadlock is DETECTED is rolled back.

Incorrect Answers

C: Only statement causing the deadlock in the session where the deadlock is detected is rolled back, not all sessions involved in the deadlock are terminated.

D: All transactions currently in progress are not terminated.

E: All database sessions do not hang indefinitely because of the deadlock: Oracle will detect and resolve deadlock situations in most cases.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 468-470
Chapter 9: Tuning Contention

QUESTION 148:

Which two statements are true of the buffer cache? (Choose two)

- A. In Oracle9i, the buffer caches can be individually sized.
- B. Blocks in the buffer cache may be in one of several states.
- C. In Oracle9i additional buffer caches can be dynamically added.
- D. The blocks in the buffer cache are managed using two or more lists, depending on the number of buffer caches in use.
- E. Blocks in the buffer cache may be in one of two states, either free or dirty.

Answer: A, B

Explanation:

In Oracle9i, the buffer caches (DEFAULT, KEEP, RECYCLE) can be individually sized. Blocks in buffer cache can be in three states: free, dirty or pinned. In Oracle9i additional buffer caches can be dynamically added by using ALTER SYSTEM command.

Incorrect Answers

C: In Oracle9i additional buffer caches can be dynamically sized, but they cannot be added by using ALTER SYSTEM command.

D: The blocks in the buffer cache are not managed using two or more lists.

E: This statement is wrong: blocks in the buffer cache may be not only in one of two states. They can be free, dirty or pinned.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 275-277
Chapter 5: Tuning the Database Buffer Cache

QUESTION 149:

You execute the following query:

```
SELECT sum(waits) * 100 / sum(gets) "Ratio",  
sum(waits) "Waits", sum(gets) "Gets"  
FROM vsrollstat;
```

If the number of waits for any class is greater than 1% of the total number of requests, what should you do?

- A. Create more rollback segments.
- B. Check for WAITS in v\$rollstat and create more extents.
- C. Increase the number of extents in the rollback segments.
- D. Check for WRAPS in v\$rollstat and create larger extents.
- E. Increase the size of the extents in the rollback segments.

Answer: A

Explanation:

You should create more rollback segments to decrease rollback segments contention.

Incorrect Answers

B: You don't need to create more extents.

C: It's not required to increase the number of extents in the rollback segments.

D: It's not possible to create larger extents: in rollback segments all extents are the same size.

E: Because of the fixed size of extents in the rollback segments you cannot increase their size.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 338-342

Chapter 7: Tuning Redo Mechanisms

QUESTION 150:

What are two valid ways to reduce lock contention? (Choose two)

- A. Use the minimum level of locking required.
- B. Use select for update statement frequently.
- C. Use lock table in exclusive mode frequently.
- D. Use the Oracle default locking level for updates.

Answer: A, D

Explanation:

There are two possible ways to reduce lock contention: use the minimum level of locking required and use the Oracle default locking level for updates.

Incorrect Answers

B: Using select for update statement frequently you will lock tables more frequently and increase lock contention by doing this.

C: Using lock table in exclusive mode frequently will also increase lock contention.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 468-470
Chapter 9: Tuning Contention

QUESTION 151:

You create the COMPANY table with the NOLOGGING attribute:

```
CREATE TABLE company (  
  company_id NUMBER(5) not null,  
  company_name VARCHAR2(4) NOT NULL,  
  address VARCHAR2(200)
```

```
]
```

```
NOLOGGING
```

```
storage (initial 321 pctincrease 0);
```

Which two operations on the COMPANY table will suppress redo entry generation?

(Choose two)

A. DELETE FROM COMPANY;

B. Direct Path loads using SQL*Loader

C. DELETE FROM COMPANY WHERE company id=1000;

D. Direct-load INSERT operation on the COMPANY table.

E. UPDATE COMPANY SET COMPANY_ADDRESS = 'One Oracle Way'

Where company_id = 100

Answer: B, D

Explanation:

Direct Path loads using SQL*Loader and direct-load INSERT operation on the COMPANY table will suppress redo entry generation.

Incorrect Answers

A: This DML command will not suppress redo entry generation.

C: This DML command will not suppress redo entry generation.

E: This DML command will not suppress redo entry generation.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 338-342
Chapter 7: Tuning Redo Mechanisms

QUESTION 152:

What are the three levels at which the Oracle optimize mode can be set? (Choose three)

A. Set at the statement level by using hints.

B. Set at the session level by using the ALTER SESSION command.

C. Set at the instance level by using the OPTIMIZER_MODE parameter.

D. Set inside the SQL code by issuing the SET OPTIMIZER command at the time of the execution.

Answer: A, B, C

Explanation:

You can set the Oracle optimize mode at the statement level, session level and instance level.

Incorrect Answers

D: It's not possible to set Oracle optimizer mode inside the SQL code.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 99-101

Chapter 3: SQL Application Tuning and Design

QUESTION 153:

Several users have called to state that their normal sub-second response time for OLTP applications is now averaging three to four seconds.

You find the following:

1. The SGA is currently at its maximum size defined by SGA_MAX_SIZE.
2. The GETHITRATIO for the library cache is 98%.
3. The reloads-to-pins is less than .25 percent when you query the V\$LIBRARYCACHE view.
4. The hit ratio in the buffer cache has fallen from its normal level of 88 to 90 percent and is only averaging 65 percent when you look at several different the reports from STATSPACK.

You have used the SHOW command to verify the current settings for the SHARED_POOL_SIZE and the DB_CACHE_SIZE parameters.

Which three steps should you take to improve response time for the OLTP application user? (choose three. First choose 1 From A and B; then choose 1 from C and D; then choose one from E and F)

A. Increase the size of the default buffer cache by using the *ALTER SYSTEM SET DB_CACHE_SIZE=nnnM;"command where nnn is a number and M stands for megabytes.

Oracle will automatically reduce the size of the shared pool to accommodate the increase in size of the buffer cache that you have specified.

B. Reduce the size of the shared pool using the "ALTER SYSTEM SET SHARED_POOL_SIZE=nnnM;"command followed by the "ALTER SYSTEM SET DB_CACHE_SIZE=nnnM;"command where the nnn; represent numbers and M stand for megabytes.

C. Query the V\$BUFFER_POOL view to verify that the buffer cache has grown to the specified size.

D. Query the V\$BUFFER_POOL_STATISTICS view to verify that the buffer cache has grown to the specified size.

E. Check the buffer cache hit ratio after the database has been running for some time. Repeat the process if the buffer cache hit ratio is less than 90%.

F. Check the buffer cache hit ratio after the database has been running for some time. Repeat the process if the buffer cache hit ratio is less than 90%, the GETHITRATIO in V\$LIBRARYCACHE is still greater than 90%, and the reloads-to-pins is still less than 1% in V\$LIBRARYCACHE.

Answer: B, C, F

Explanation:

These steps are necessary to optimize response time for the OLTP application users.

Incorrect Answers

A: Oracle will not automatically reduce the size of the shared pool to accommodate the increase in size of the buffer cache that you have specified.

D: You should query the V\$BUFFER_POOL view, not V\$BUFFER_POOL_STATISTICS view because V\$BUFFER_POOL view contains information about current DB_CACHE_SIZE.

E: You need to continue process if the GETHITRATIO in V\$LIBRARYCACHE is still greater than 90%, and the reloads-to-pins is still less than 1% in V\$LIBRARYCACHE because you don't want to decrease library cache performance by tuning buffer cache.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 205-208

Chapter 4: Tuning the Shared Pool

QUESTION 154:

Last week when you created your Oracle9i database, you set UNDO_MANAGEMENT to AUTO in the parameter file, and you specified a directory for the parameter DB_CREATE_FILE_DEST. You attempt to create an UNDO TABLESPACE named MY_UNDO. You do not use a data file subclause. Which statement is true?

- A. The creation of MY_UNDO fails.
- B. Oracle places the file in \$ORACLE_HOME/oradata/<sid> directory.
- C. Oracle creates a data file and placed it in \$ORACLE_HOME/dbs directory.
- D. Oracle creates a data file in the directory specified in DB_CREATE_FILE_DEST.

Answer: D

Explanation:

Oracle will create a data file in the directory specified in DB_CREATE_FILE_DEST if you set UNDO_MANAGEMENT to AUTO in the parameter file, and you specified a directory for the parameter DB_CREATE_FILE_DEST.

Incorrect Answers

A: The creation of MY_UNDO does not fail.

B: Oracle will not place this file in \$ORACLE_HOME/oradata/<sid> directory if you specified a directory for the parameter DB_CREATE_FILE_DEST.

C: Oracle will not place this file in \$ORACLE_HOME/dbs directory if you specified a

directory for the parameter DB_CREATE_FILE_DEST.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 445-446

Chapter 8: Tuning Disk I/O

QUESTION 155:

Which statement is true about areas in the shared pool?

- A. Invalidations in the library cache are caused by misses in the data dictionary cache.
 - B. Using bind variables rather than constants and as much generic code as possible improves the performance of library cache.
 - C. If the ratio of GETMISSES to GETS found in the V\$ROWCACHE is less than 15 percent you know that both the data dictionary cache and the library cache in the shared pool are tuned correctly.
 - D. The init.ora parameter SESSION_CACHED_CURSOR sets the number of execution plans cached in the shared pool area.
- You can view the actual execution plan used by querying the V\$SQL_PLAN view.

Answer: B

Explanation:

You can improve the performance of library cache by using bind variables rather than constants and as much generic code as possible.

Incorrect Answers

A: Invalidations in the library cache are not caused by misses in the data dictionary cache.

C: This statement does not prove that the library cache in the shared pool is tuned properly.

D: To enable caching of session cursors, you must set the initialization parameter SESSION_CACHED_CURSORS. The value of this parameter is a positive integer specifying the maximum number of session cursors kept in the cache. An LRU algorithm removes entries in the session cursor cache to make room for new entries when needed.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 205-208

Chapter 4: Tuning the Shared Pool

QUESTION 156:

You run this script and discover that the value of request failures is more than zero:

```
SELECT * from v$shared_pool_reserved  
WHERE request_failures>0;
```

Which two init.ora parameters would you change to reduce REQUEST_FAILURES to 0?
(Choose two)

- A. MTS_SERVERS
- B. DB_BLOCK_SIZE
- C. SORT_AREA_SIZE

- D. SHARED_POOL_SIZE
- E. DB_BLOCK_BUFFERS
- F. SHARED_POOL_RESERVED_SIZE

Answer: D, F

Explanation:

The reserved pool is too small when the value for REQUEST_FAILURES is more than zero and increasing. To resolve this, increase the value for the SHARED_POOL_RESERVED_SIZE and SHARED_POOL_SIZE accordingly. The settings you select for these parameters depend on your system's SGA size constraints. Increasing the value of SHARED_POOL_RESERVED_SIZE increases the amount of memory available on the reserved list without having an effect on users who do not allocate memory from the reserved list. If you have enabled the reserved list, then decrease the value for SHARED_POOL_RESERVED_SIZE. If you have not enabled the reserved list, then you could increase SHARED_POOL_SIZE.

Incorrect Answers

A: MTS_SERVERS parameter is not related with statistics in V\$SHARED_POOL_RESERVED view.

B: Changes in DB_BLOCK_SIZE will not affect SHARED_POOL_RESERVED_SIZE value.

C: Statistics in V\$SHARED_POOL_RESERVED cannot be improved by SORT_AREA_SIZE parameter changes.

E: DB_BLOCK_BUFFERS parameter is deprecated in favor of the DB_CACHE_SIZE parameter. It specifies the number of database buffers in the buffer cache. It is one of several parameters that contribute to the total memory requirements of the SGA of an instance.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 205-208
Chapter 4: Tuning the Shared Pool

QUESTION 157:

Which DBMS package's procedures can you use to copy statistics from an Oracle9i production database to a test database, in order to facilitate tuning?

- A. DBMS_STATS
- B. DBMS_ANALYZE
- C. DBMS_EXPORT_EXTENSION
- D. DBMS_RMGR_PLAN_EXPORT
- E. DBMS_AQ_EXP_HISTORY_TABLES

Answer: A

Explanation:

The PL/SQL package DBMS_STATS lets you generate and manage statistics for cost-based optimization. You can use this package to gather, modify, view, export,

import, and delete statistics. You can also use this package to identify or name statistics gathered. The IMPORT_* procedures retrieve statistics, including user-defined statistics, from the STATTAB table and store them in the dictionary. Because the SET*_STATS and GET*_STATS interfaces are supported for user-defined statistics, user-defined statistics can be copied to another database using this interface.

Incorrect Answers

B: There is no DBMS_ANALYZE DBMS package in Oracle9i.

C: The DBMS_EXPORT_EXTENSION DBMS package implements PL/SQL extensions to Export in Oracle9i.

D: The DBMS_RMGR_PLAN_EXPORT DBMS package in Oracle9i is not used to copy statistics.

E: The DBMS_AQ_EXP_HISTORY_TABLES DBMS package in Oracle9i is not used to copy statistics.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 40-41
Chapter 2: Sources of Tuning Information

QUESTION 158:

Which package would you use to grant resource groups to users?

- A. DBMS_REPCAT_AUTH
- B. DBMS_SPACE_ADMIN
- C. DBMS_REPCAT_ADMIN
- D. DBMS_RESOURCE_MANAGER
- E. DBMS_RESOURCE_MANAGER_PRIVS

Answer: E

Explanation:

The DBMS_RESOURCE_MANAGER_PRIVS package maintains privileges associated with the Resource Manager.

Incorrect Answers

A: There is no DBMS_REPCAT_AUTH DBMS package in Oracle 9i.

B: The DBMS_SPACE_ADMIN package provides functionality for locally managed tablespaces.

C: DBMS_REPCAT_ADMIN enables you to create users with the privileges needed by the symmetric replication facility.

D: The DBMS_RESOURCE_MANAGER package maintains plans, consumer groups, and plan directives. It also provides semantics so that you may group together changes to the plan schema.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 517-547
Chapter 10: Operating System Tuning

QUESTION 159:

1Z0-033

Which two statements are true regarding the use of STATSPACK? (Choose two)

- A. Since undo segments are new as of Oracle9i, STATSPACK does not cover undo segments in the rollback segment statistics section.
- B. If snapshot collection is automated through DBMS_JOB package, the interval of statistics collection cannot be changed.
- C. Unlike UTLBSTAT/UTLESTAT, STATSPACK allows the flexibility of basing a report on any data points selected over the same uninterrupted running timeframe of each instance.
- D. Each time the 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.
- E. Since many init<SID>.ora parameters are dynamic, STATSPACK guarantees that the current parameter values shown in the output will be the ones used if the database is restarted.
- F. STATSPACK.snap procedure will store the current values for the performance statistics in the STATSPACK tables, which can be used as a baseline snapshot for comparison with another snapshot taken at a later time.

Answer: C, F

Explanation:

STATSPACK guarantees that the current parameter values shown in the output will be the ones used if the database is restarted

Incorrect Answers

- A: STATSPACK supports undo segments in the rollback segment statistics section.
- B: The interval of statistics collection can be changed.
- D: STATSPACK will not automatically generate a report that displays the statistical output. Statspack separates data collection from report generation. Data is collected when a snapshot is taken. The performance engineer then runs the performance report and views the data collected.
- E: STATSPACK guarantees that the current parameter values shown in the output will be the ones used. The database does not require to be restarted.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 40-41

Chapter 2: Sources of Tuning Information

QUESTION 160:

It is possible for a DBA or Application Developer to override Oracle's default locking. In this case, the locking procedures set by the DBA or Application Developer must satisfy which two criteria? (Choose two)

- A. Ensure that data locks are not held.
- B. Ensure that implicit commits are used.
- C. Ensure that the Savepoint clause is used.
- D. Ensure that data concurrency is acceptable.

E. Ensure that all locks are held for the shortest time possible.

Answer: D, E

Explanation:

By overriding Oracle's default locking you need to be sure that data concurrency is acceptable and that all locks are held for the shortest time possible.

Incorrect Answers

A: Data locks should be held at least for short time.

B: It is not required to use implicit commits.

C: It is not required to use Savepoint clause.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 468-470

Chapter 9: Tuning Contention

QUESTION 161:

You need some extra space in your database, so you issue this command:

```
ALTER TABLE my_table DEALLOCATE UNUSED;
```

What happens?

A. All blocks of MY_TABLE that contain no rows are de-allocated.

B. All blocks below the high-water mark of MY_TABLE are deallocated.

C. The high-water mark of MY_TABLE is recalculated and stored in the segment header.

D. All blocks above the high-water mark of MY_TABLE are deallocated, as long as the high-water mark is beyond the minimum creation size of the segment.

Answer: D

Explanation:

If the high-water mark is beyond the minimum creation size of the segment, all blocks above the high-water mark of MY_TABLE are deallocated.

Incorrect Answers

A: All blocks above the high-water mark of MY_TABLE are deallocated, not all blocks containing no rows.

B: All blocks ABOVE, not below, the high-water mark of MY_TABLE are deallocated.

C: The high-water mark of MY_TABLE is not recalculated.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 445-446

Chapter 8: Tuning Disk I/O

QUESTION 162:

In which two ways can you reduce the amount of sorting that is performed? (Choose two)

A. Use UNION instead of UNION ALL.

B. Use NOSORT when creating tables.

C. Use NOSORT when creating indexes.

- D. Use COMPUTE instead of ESTIMATE when analyzing objects.
- E. Reduce the number of users that have the sort privilege.
- F. Create appropriate indexes on tables that are joined often.

Answer: C, F

Explanation:

One cause of sorting is the creation of indexes. Creating an index for a table involves sorting all rows in the table based on the values of the indexed columns. However, Oracle lets you create indexes without sorting. If the rows in the table are loaded in ascending order, then you can create the index faster without sorting. Another way to reduce the amount of sorting is to create appropriate indexes on tables that are joined often.

Incorrect Answers

- A: Usage of UNION command will still require sorting.
 - B: The NOSORT clause is used with index creation, not table.
 - D: Using ESTIMATE, not COMPUTE, clause you can reduce the amount of sorting.
 - E: There are no "sorting" privileges in Oracle.
- OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 445-446
Chapter 8: Tuning Disk I/O

QUESTION 163:

You want to monitor row migrations but you don't want to create the CHAINED_ROWS table manually. Which script must you run?

- A. CATPROC
- B. CATPERF
- C. UTLBSTAT
- D. UTLCHAIN

Answer: D

Explanation:

The definition of a sample output table named CHAINED_ROWS appears in a SQL script available on your distribution medium. The common name of this script is UTLCHN1.SQL, although its exact name and location varies depending on your platform. Your output table must have the same column names, datatypes, and sizes as the CHAINED_ROWS table.

Incorrect Answers

- A: CATPROC script does not create this table. It is required to load the initial PL/SQL environment.
 - B: There is no CATPERF in Oracle9i.
 - C: UTLBSTAT/UTLESTAT scripts are used for statistic collection.
- OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 40-41
Chapter 2: Sources of Tuning Information

QUESTION 164:

Which two types of information found in the alert log are relevant for database performance tuning? (Choose two)

- A. database name
- B. internal errors
- C. controlfile reuse
- D. core dump destination
- E. checkpoint not complete

Answer: B, E

Explanation:

Internal errors and not completion of checkpoint can be relevant for database performance tuning.

Incorrect Answers

A: Database name has no effect on database performance.

C: Controlfile reuse is not related with performance tuning.

D: Core dump destination has nothing to do with database performance.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 40-41

Chapter 2: Sources of Tuning Information

QUESTION 165:

Which three statements are true about stored outlines in Oracle9i? (Choose three)

- A. Outlines always use the rule-based optimizer.
- B. The USE_STORED_OUTLINES parameter can be set at the session level.
- C. The same SQL statement can have a stored outline in more than one category.
- D. If USE_STORED_OUTLINES is set to TRUE, then outlines from the DEFAULT category are used.

Answer: B, C, D

Explanation:

The USE_STORED_OUTLINES parameter can be set at the session level. Outlines can be categorized to simplify the management task. The CREATE OUTLINE statement allows for specification of a category. The DEFAULT category is chosen if unspecified. Likewise, the CREATE_STORED_OUTLINES parameter lets you specify a category name, where specifying true produces outlines in the DEFAULT category.

Incorrect Answers

A: By creating outlines for an application before switching to cost-based optimization, the plans generated by the rule-based optimizer can be used, while statements generated by newly written applications developed after the switch use cost-based plans.

QUESTION 166:

When does the Oracle server put index segment blocks back on the free list?

- A. When you delete 25% of the entries from an index block.
- B. When you delete 50% of the entries from an index block.
- C. When you delete 75% of the entries from an index block.
- D. When you delete 100% of the entries from an index block.

Answer: D

Explanation:

The Oracle server put index segment blocks back on the free list when you delete 100% of the entries from an index block.

Incorrect Answers

A: The Oracle server put index segment blocks back on the free list when you delete 100%, not 25%, of the entries from an index block.

B: The Oracle server put index segment blocks back on the free list when you delete 100%, not 50%, of the entries from an index block.

C: The Oracle server put index segment blocks back on the free list when you delete 100%, not 75%, of the entries from an index block.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 445-446
Chapter 8: Tuning Disk I/O

QUESTION 167:

Which is a correct guideline for resolving wait times on the request queue when using Oracle Shared Servers?

- A. Reduce the number of dispatchers, thus reducing the requests to a number that is manageable by the shared servers.
- B. Increase the size of the large pool, so that less time is spent performing I/O connected with the session data.
- C. Increase the size of the shared pool, so that less time is spent performing I/O connected with the session data.
- D. If the value of `SERVERS_HIGHWATER` reaches the value set for the `MAX_SHARED_SERVERS` initialization parameter, consider raising the value of `SHARED_SERVERS`.

Answer: D

Explanation:

Make sure that the optional `MAX_SHARED_SERVERS` line is correctly set. For

example: `MAX_SHARED_SERVERS = 5`. This line sets the upper bound on the total number of shared servers PMON can create, based on the peak load of the system. This should be set high enough so that all requests can be serviced, but not so high that the system swaps if they are reached. The purpose of this parameter is to prevent the server from swapping. Run the following script to see what the highwater mark is for the number of servers running, and then set `MAX_SHARED_SERVERS` to more than this.

```
SELECT maximum_connections "MAX CONN", servers_started  
"STARTED", servers_terminated "TERMINATED", servers_highwater  
"HIGHWATER" FROM V$SHARED_SERVER_MONOTER;
```

MAKE sure that the optional `SHARED_SERVERS` line is correctly set. For example:
`SHARED_SERVERS = 5`. This is the total number of shared servers started when the database is started. It also represents the total number of shared servers PMON tries to keep. It should be the total number of servers expected to be used when the database is active. `MAX_SHARED_SERVERS` is intended to handle peak load.

Incorrect Answers

A: By reducing the number of dispatcher you can decrease overall database performance.

B: Increasing the size of the large pool will not solve the problem.

C: Increasing the size of the shared pool will not solve the problem.

OCP: Oracle 9i Performance Tuning Study Guide, Jaseph C. Johnson, p. 517-547

Chapter 10: Operating System Tuning

QUESTION 168:

Which command can be used to turn on the buffer cache advisory feature?

A. `ALTER SESSION SET DB_CACHE_ADVICE=on;`

B. `ALTER SYSTEM SET DB_CACHE_ADVICE=on;`

C. `ALTER DATABASE SET DB_CACHE_ADVICE=on;`

D. `ALTER DATABASE BUFFER POOL SET DB_CACHE_ADVICE=on;`

Answer: B

Explanation:

`DB_CACHE_ADVICE` enables or disables statistics gathering used for predicting behavior with different cache sizes through the `V$DB_CACHE_ADVICE` performance view. It's dynamic parameter class: it uses the `ALTER SYSTEM` command.

Incorrect Answers

A: To turn on the buffer cache advisory feature you need to use the `ALTER SYSTEM` command.

C: To turn on the buffer cache advisory feature you need to use the `ALTER SYSTEM` command.

D: To turn on the buffer cache advisory feature you need to use the `ALTER SYSTEM` command.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 275-277

Chapter 5: Tuning the Database Buffer Cache

QUESTION 169:

You notice row migration in the database and want to prevent it from continuing. What can be done to avoid future row migration?

- A. Recreate the database with smaller DB_BLOCK_SIZE.
- B. Increase the PCTFREE setting on the tables with row migration problems.
- C. Increase the MAXEXTENTS setting on the tables with row migration problems.
- D. Increase the NEXT setting for the tables with row migration problems so that the size of the next extent allocated is bigger.

Answer: B

Explanation:

If an UPDATE statement increases the amount of data in a row so that the row no longer fits in its data blocks, the Oracle Server tries to find another block with enough free space to hold the entire row. If such a block is available, the Oracle Server migrates the data for the entire row to the new data. A forwarding address is left at the original row's block.

This is called migrating a row; the rowid of a migrated row does not change.

The effect of row fragmentation on performance is that each time a migrated row is needed for a query, an additional I/O call is required for the row piece that is not stored in the original block.

It is possible to minimize unnecessary row fragmentation for a table by correctly setting the table's PCTFREE value. Being too pessimistic on PCTFREE can waste some space but will more or less guarantee no row migration while being too optimistic can result in row migration and reduced performance. Instead of trying to save every byte in a block, Oracle recommends erring on the side of setting PCTFREE high rather than low. In most cases this will be a difference of only 5%. For tables that do not have varying width columns and/or do not get updated, this parameter can be safely ignored or set to a low value.

Incorrect Answers

A: Smaller DB_BLOCK_SIZE will increase of row migration possibility.

C: The MAXEXTENTS setting on the tables has nothing to do with row migration problems.

D: If you increase the NEXT setting for the tables with row migration problems so that the size of the next extent allocated is bigger. But row migration happened with UPDATE statements affecting CURRENT blocks.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 445-446

Chapter 8: Tuning Disk I/O

QUESTION 170:

What is the correct syntax to ensure an object is not aged out of the shared pool?

- A. SQL> EXECUTE dbms_shared_pool.pin('package_name');
- B. SQL> EXECUTE dbms_shared_pool.keep('package_name');

1Z0-033

- C. SQL> EXECUTE dbms_shared_pool.unkeep('package_name');
- D. SQL> EXECUTE dbms_shared_pool.retain('package_name');
- E. SQL> EXECUTE dbms_shared_pool.refresh('package_name');

Answer: B

Explanation:

To use the DBMS_SHARED_POOL package to pin a SQL or PL/SQL area, complete the following steps.

1. Decide which packages or cursors to pin in memory.
2. Start up the database.
3. Make the call to DBMS_SHARED_POOL.KEEP to pin your objects.

Incorrect Answers

- A: There is no PIN procedure in the DBMS_SHARED_POOL package.
 - C: This procedure unkeeps the named object.
 - D: There is no RETAIL procedure in the DBMS_SHARED_POOL package.
 - E: There is no REFRESH procedure in the DBMS_SHARED_POOL package.
- OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 205-208
Chapter 4: Tuning the Shared Pool

QUESTION 171:

Which statement requires an exclusive DDL lock?

- A. AUDIT statement
- B. SELECT statement
- C. DROP TABLE statement
- D. CREATE PROCEDURE statement

Answer: C

Explanation:

DROP TABLE statement requires an exclusive DDL lock.

Incorrect Answers

- A: AUDIT statement does not require an exclusive DDL lock.
 - B: SELECT statement does not require any DDL lock.
 - D: This statement does not have requirement for an exclusive DDL lock.
- OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 468-470
Chapter 9: Tuning Contention

QUESTION 172:

Which three actions can help avoid the disadvantages of dynamic extent allocation?
(Choose three)

- A. Use Oracle default sizing.

- B. Size the segments appropriately.
- C. Use globally managed tablespaces.
- D. Use locally managed tablespaces
- E. Use recursive space management operations.
- F. Monitor segments ready to extend and extend them manually.

Answer: B, D, F

Explanation:

To avoid the disadvantages of dynamic extent allocation size the segments appropriately, use locally managed tablespaces and monitor segments ready to extend and extend them manually.

Incorrect Answers

A: Usage of Oracle default settings will not help to avoid the disadvantages of dynamic extent allocation.

C: There are no globally managed tablespaces in Oracle. Only dictionary and locally managed tablespaces exist.

E: Using recursive space management operations will not help to avoid dynamic extent allocation.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 445-446
Chapter 8: Tuning Disk I/O

QUESTION 173:

When are temporary segments created?

- A. Upon startup of the database.
- B. At the time of the first sort.
- C. When the TEMPORARY tablespace is created.
- D. When the first sort that requires disk space occurs.
- E. When the "CREATE SORT SEGMENT..." statement is run.

Answer: D

Explanation:

Within a temporary tablespace, all sort operations for a given instance and tablespace share a single sort segment. Sort segments exist for every instance that performs sort operations within a given tablespace. The sort segment is created by the first statement that uses a temporary tablespace for sorting, after startup, and is released only at shutdown. An extent cannot be shared by multiple transactions.

Incorrect Answers

A: They are created by the first statement that uses a temporary tablespace for sorting, AFTER startup of the database.

B: They are created not only at the time of the first sort. Sort should use a temporary tablespace for sorting.

C: They are created by the first statement that uses a temporary tablespace for sorting,

not when the TEMPORARY tablespace is created.

E: There is no CREATE SORT SEGMENT in Oracle.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 445-446

Chapter 8: Tuning Disk I/O

QUESTION 174:

CKPROF is _____.

- A. A utility that enables SQL code to run faster.
- B. Used to format the trace file generated by SQL_TRACE
- C. A diagnostic tool used to capture operating system statistics.
- D. A utility used to capture the statistics of all SQL currently running in the database.

Answer: B

Explanation:

CKPROF is used to format the trace file generated by SQL_TRACE.

You can run the CKPROF program to format the contents of the trace file and place the output into a readable output file. Optionally, CKPROF can also:

1. Determine the execution plans of SQL statements.
2. Create a SQL script that stores the statistics in the database.

CKPROF reports each statement executed with the resources it has consumed, the number of times it was called, and the number of rows which it processed. This information lets you easily locate those statements that are using the greatest resource. With experience or with baselines available, you can assess whether the resources used are reasonable given the work done.

Incorrect Answers

A: It does not enable SQL code to run faster.

C: This is not a diagnostic tool used to capture operating system statistics.

D: CKPROF is not used to capture the statistics of all SQL currently running in the database.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 40-41

Chapter 2: Sources of Tuning Information

QUESTION 175:

What are the three most common problems encountered while trying to tune a database?
(Choose three)

- A. Poor application design that does not make good use of indexes.
- B. Applications not written to use bind variable to take advantage of shared SQL statements.
- C. Poor application design that require too many tables to be joined together to resolve queries.
- D. Applications not closing cursors in a timely fashion because the init.ora parameter

OPEN_CURSORS was set too high.

Answer: A, B, D

Explanation:

All these reasons can explain poor database performance.

Incorrect Answers

C: Number of tables that need to be joined should not affect performance in good tuned database.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 99-101
Chapter 3: SQL Application Tuning and Design

QUESTION 176:

Which three can you use to monitor locks? (Choose three)

- A. V\$LOCK
- B. DBA_WAITERS
- C. V\$LOCKED_OBJECT
- D. DBA_PENDING_TRANSACTIONS

Answer: A, B, C

Explanation:

All these views can be used to monitor locks. DBA_WAITERS shows all the sessions that are waiting for a lock, but do not hold locks for which another session is waiting. V\$LOCKED_OBJECT view lists all locks acquired by every transaction on the system.

Incorrect Answers

D: DBA_PENDING_TRANSACTIONS describes unresolved transactions (either due to failure or if the coordinator has not sent a commit/rollback).

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 468-470
Chapter 9: Tuning Contention

QUESTION 177:

Which two methods can you use to calculate the size of the recycle pool? (Choose two)

- A. Total size the buffer cache blocks used by the object.
- B. Use the direct reads statistic from a tracing tool.
- C. Use the physical reads statistic from a tracing tool.
- D. Use the indirect reads statistic from a tracing tool.

Answer: A, C

Explanation:

To calculate the size of the recycle pool you can use total size the buffer cache blocks

used by the object and the physical reads statistic from the tracing tool.

Incorrect Answers

B: The direct reads statistic from a tracing tool cannot be used for this purpose because of nature of the recycling pool.

D: The indirect reads statistic from a tracing tool will not help you to calculate the size of the recycle pool.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 205-208

Chapter 4: Tuning the Shared Pool

QUESTION 178:

You want to find the high water mark and the number of blocks above the high water mark without analyzing the table. Which method will provide this information?

- A. Use DBMS_OUTPUT package.
- B. Use the ANALYZE command.
- C. Use DBMS_SQL.PARSE procedure.
- D. Use DBMS_SPACE.UNUSED_SPACE procedure.

Answer: D

Explanation:

The DBMS_SPACE.UNUSED_SPACE procedure will help you to find the high water mark and the number of blocks above the high water mark without analyzing the table. This procedure returns information about unused space in an object (table, index, or cluster).

Incorrect Answers

A: Package DBMS_OUTPUT enables you to display output from PL/SQL blocks and subprograms, which makes it easier to test and debug them.

B: The task is not to use the ANALYZE command.

C: The PARSE procedure from the DBMS_SQL package can be used to parse DML or DDL statement.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 40-41

Chapter 2: Sources of Tuning Information

QUESTION 179:

What may happen if you size your rollback segments too small?

- A. The long transaction switches to SYSROL rollback.
- B. The long transaction fails for lack of rollback space.
- C. The long transaction hangs until the DBA alters MINEXTENTS.
- D. The long transaction is forced to switch to another rollback segment.

Answer: B

Explanation:

The long transaction can fail for lack of rollback space if you size your rollback segments too small.

Incorrect Answers

A: Transaction cannot be switched to another rollback segment.

C: MINEXTENTS parameter has nothing to do with rollback segments. The long transaction will just fail.

D: Transaction cannot be switched to another rollback segment.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 338-342

Chapter 7: Tuning Redo Mechanisms

QUESTION 180:

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

- A. The database is experiencing database-intensive work.
- B. The database is primary 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 primary used for an interactive application where dedicated servers are mainly idle.

Answer: D, F

Explanation:

Oracle Shared Server configuration will be very helpful when an OLTP application is running on a machine approaching memory resource limits or the database is primary used for an interactive application where dedicated servers are mainly idle. The idle server process is holding system resources. The shared server architecture eliminates the need for a dedicated server process for each connection

Incorrect Answers

A: If the database is experiencing database-intensive work all dedicated servers are busy. But Oracle Shared Server architecture is more useful when dedicated servers are mainly idle.

B: The Oracle Shared Server configuration is more useful for OLTP systems with many DML operations, not for DSS environment with long-running transactions.

C: Oracle Shared Server will be not the best in this situation.

E: It is required when you NEED to scale up the number of concurrent connections to the database.

OCP: Oracle 9i Performance Tuning Study Guide, Jaseph C. Johnson, p. 517-547

Chapter 10: Operating System Tuning

QUESTION 181:

Which four activities cause a sort? (Choose four)

- A. Creation of an index.
- B. Execution of the ANALYZE command.
- C. Use of the ORDER BY and GROUP BY clauses.
- D. Use of insert operations involving index maintenance.
- E. Use of the UNION, INTERSECT and MINUS operators.

Answer: A, C, D, E

Explanation:

All these activities can cause a sort operation. Only if you are creating an index with NOSORT option sorting will not have place.

Incorrect Answers

B: Execution of the ANALYZE command does not require sorting.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 445-446

Chapter 8: Tuning Disk I/O

QUESTION 182:

The Accounting department has contacted you concerning their perceived poor performance of the new accounting package that was installed last week. Which three actions would help diagnose the problem? (Choose three)

- A. Query the DBA_ROLLBACK_SEGS table to see if there is contention for rollback segment headers while the accounting application is being run.
- B. Query DBA_DATA_FILES while the accounting application is being used to determine if any of the data files are being used excessively.
- C. Query the V\$SQL_PLAN table to determine which, if any, indexes are being used for the accounting tables.
- D. Query the V\$FILESTAT table while the accounting application is being used to determine if any of the data files are being used excessively.
- E. Query the V\$LOGFILE table to see if an excessive amount of redo is being generated while the accounting application is being run.
- F. Generate performance monitoring reports using UTLBSTAT and UTLESTAT or the STATSPACK while the account application is being used and compare them to your previously generated base line reports from the same tool.

Answer: C, D, F

Explanation:

The V\$SQL_PLAN view provides a way of examining the execution plan for cursors that were executed and are still cached. The V\$FILESTAT view contains detailed file I/O statistics for each file, including the number of I/Os for each file and the average read time. Generate performance monitoring reports using UTLBSTAT and UTLESTAT or the STATSPACK will also help to find the reason of poor performance.

Incorrect Answers

A: The DBA_ROLLBACK_SEGS data dictionary view stores information about the rollback segments of a database. But this information cannot be used to check contention for rollback segment headers.

B: The DBA_DATA_FILES data dictionary view provides descriptive information about each datafile, including the tablespace to which it belongs and the file id. The file id can be used to join with other views for detail information, but itself it will not give you information if any of the data files are being used excessively.

E: This view just identifies redo log groups and members and member status.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 40-41
Chapter 2: Sources of Tuning Information

QUESTION 183:

Which three are used when administering systems with the database Resource Manager?
(Choose three)

- A. Histograms
- B. Explain Plan
- C. Resource Plan
- D. Resource User Group
- E. Resource Consumer Group
- F. Resource Plan Directives
- G. Resource Plan Assignments

Answer: C, E, F

Explanation:

Resource Plan, Resource Consumer Group and Resource Plan Directives are used when administering systems with the database Resource Manager.

Incorrect Answers

A: Histograms are not used with the database Resource Manager.

B: Explain Plan has no relation with the database Resource Manager.

D: There is Resource Consumer Group, not Resource User Group in Resource Manager.

G: There is no Resource Plan Assignments for Resource Manager.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 517-547
Chapter 10: Operating System Tuning

QUESTION 184:

There are two users, John and Susan, who are updating the EMPLOYEE table in the following order:

1. First, John issues this SQL update:

```
UPDATE EMPLOYEE SET LAST_NAME='SMITH'where ID =200;
```

2. Next, Susan issues this SQL update:

```
UPDATE EMPLOYEE SET SALARY=50000 WHERE ID=250;
```

3. Next, John issues this SQL update:

1Z0-033

UPDATE EMPLOYEE SET LAST_NAME'BAKER'WHERE ID =250;

4. FINALLY, Susan issues this SQL update:

UPDATE EMPLOYEE SET SALARY=60000 WHERE ID=200;

What will be the result?

- A. Oracle kills Susan's session to prevent a deadlock.
- B. Oracle will detect a deadlock and roll back Susan's session.
- C. Oracle kills both John's and Susan's statements to prevent a deadlock.
- D. Both John's and Susan's sessions would hang indefinitely because of a deadlock.
- E. Oracle will detect a deadlock and roll back the statement causing the deadlock.

Answer: E

Explanation:

Oracle will detect a deadlock and roll back the statement causing the deadlock.

Incorrect Answers

A: Oracle will detect a deadlock FIRST and roll back the transaction detected a deadlock.

It will not kill any session.

B: Oracle will detect a deadlock and roll back the statement causing the deadlock. It can be transaction from Susan's session or from John's session.

C: Oracle will not kill both statements, it will only rollback one, detecting a deadlock.

D: Both John's and Susan's sessions would not hang indefinitely: Oracle will handle a deadlock in most cases.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 468-470

Chapter 9: Tuning Contention

QUESTION 185:

Which two statements are valid settings for the initialization parameter 'USE_STORED_OUTLINES' ? (Choose two)

- A. SQL>alter session set USE_STORED_OUTLINES=TRUE;
- B. SQL>alter session set USE_STORED_OUTLINES=PUBLIC;
- C. SQL>alter session set USE_STORED_OUTLINES=PRIVATE;
- C. SQL>alter session set USE_STORED_OUTLINES=DEFAULT;
- E. SQL> alter session set USE_STORED_OUTLINES = <outline name>;
- F. SQL> alter session set USE_STORED_OUTLINES = <category name>;

Answer: A, F

Explanation:

There are two valid settings for the initialization parameter USE_STORED_OUTLINES: TRUE or category name.

Incorrect Answers

B: This statement is incorrect.

C: This statement is incorrect.

D: This statement is incorrect.

E: This initialization can use category name, not outline name.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 99-101

Chapter 3: SQL Application Tuning and Design

QUESTION 186:

Which two statements are true regarding the use of DB_CACHE_ADVICE init.ora parameter? (Choose two)

A. Setting the parameter to READY reserves space in the buffer cache to store information about different buffer cache sizes, but no CPU overhead is incurred.

B. Setting the parameter to READY reserves space in the shared pool to store information about different buffer cache sizes, but no CPU overhead is incurred.

C. Setting the parameter to ON reserves space in the buffer cache to store information about different buffer cache sizes, and CPU overhead is incurred as statistics are collected.

D. The V\$DB_CACHE_ADVICE view contains information that predicts the estimated number of physical reads for different cache sizes for each buffer cache setup in the SGA.

Answer: C, D

Explanation:

Setting the parameter to READY makes sure that advisory is turned off but the memory for the advisory remains allocated. Allocating the memory before the advisory is actually turned on avoids the risk

of an error when you switch the parameter to ON. Each cache size simulated has its own row in this view, with the predicted physical I/O activity that would take place for that size. The DB_CACHE_ADVICE parameter is dynamic, so the advisory can be enabled and disabled dynamically to allow you to collect advisory data for a specific workload.

Two minor overheads are associated with this advisory:

1. CPU: When the advisory is enabled, there is a small increase in CPU usage, because additional bookkeeping is required.

2. Memory: The advisory requires memory to be allocated from the shared pool (about 100 bytes for each buffer).

The V\$DB_CACHE_ADVICE view is populated when the DB_CACHE_ADVICE parameter is set to ON. This view shows the simulated miss rates for a range of potential buffer cache sizes.

Incorrect Answers

A: Setting the parameter to ON, not to READY, reserves space in the buffer cache to store information about different buffer cache sizes

B: Setting the parameter to READY reserves space in the buffer cache. It has nothing to do with the shared pool.

QUESTION 187:

Which three statements are true with respect to has clusters? (Choose three)

- A. Full table scans are generally faster on cluster tables than on non-clustered tables.
- B. Hash clusters may be desirable when the number of key values is predictable and key values are evenly distributes.
- C. If tables in a hash cluster require more space than the initial allocation for the cluster, performance degradation can be substantial because overflow blocks are required.
- D. Storing a single table in a hash cluster can be useful regardless of whether the table is joined frequently with other tables or not, provided other criteria for choosing a hash cluster are met.

Answer: B, C, D

Explanation:

These statements are true with respect to clusters.

Incorrect Answers

A: Full table scans are generally faster on non-cluster tables than on clustered tables.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 445-446
Chapter 8: Tuning Disk I/O

QUESTION 188:

You are considering Operating System striping for your tablespaces. Which two statements are true? (Choose two)

- A. The stripe is dependent on the Oracle block size.
- B. Controller based striping will usually perform better than O/S based striping.
- C. O/S based striping will usually perform better than controller bases striping.
- D. You should choose a stripe size that will give you one physical read for each Oracle server read.

Answer: B, D

Explanation:

Controller based striping will usually perform better than O/S based striping: more disks with independent controllers will give you more effectiveness than O/S based striping.

You should choose a stripe size that will give you one physical read for each Oracle server read.

Incorrect Answers

A: The stripe is independent on the Oracle block size.

C: Controller based striping will usually perform better than O/S based striping.

QUESTION 189:

While you were on vacation, several new applications were added to a database that you manage. Another DBA did his best to tune the database while you were gone, and increased the buffer cache hit ratio from 85 percent to 94 percent after increasing the size of the `DB_CACHE_SIZE` from 96 MB to 128 MB.

As a result, the size of the SGA is now 256 MB, which is the maximum amount of memory that can be allocated for this database. You have set the `init.ora` parameter `SGA_MAX_SIZE = 256 MB` to meet that restriction. In addition, the `init.ora` parameter `SHARED_POOL_SIZE` is set to 48 MB.

Your problem is that the OLTP applications are not performing with their normal 3-second response time as they did in the past.

What can you do to try to get the response time for OLTP applications to less than three seconds?

- A. Check the buffer cache hit ratio to make sure that it is greater than 95 percent. If it is not greater than 95 percent, decrease the `SHARED_POOL_SIZE` to 32 MB and allocate the remaining 16 MB to the buffer cache by setting the `DB_CACHE_SIZE` parameter to 144 MB.
- B. Check that the percentage of misses in the entire data dictionary cache is less than two percent by querying the `V$ROWCACHE` view. If it is not less than two percent, then decrease the `DB_CACHE_SIZE` to 112 MB and allocate the remaining 16 MB to the shared pool by setting the `init.ora` parameter `SHARED_POOL_SIZE` to 64 MB.
- C. Query the `V$LIBRARYCACHE` view to obtain the number of `RELOADS` and `PINS` that have taken place. If the number of reloads is greater than 1 percent of the pins, de-allocate 16 MB from the buffer cache by setting the `DB_CACHE_SIZE` parameter to 112 MB and add the 16 MB to the shared pool by setting the `init.ora` parameter `SHARED_POOL_SIZE` to 64 MB.
- D. Query the `V$LIBRARYCACHE` view and look at the `GETHITRATIO` column. If the `GETHITRATIO` is less than 99 percent for the `SQL AREA` namespace, de-allocate 16 MB from the buffer cache by setting the `DB_CACHE_SIZE` parameter to 112 MB and add the 16 MB to the shared pool by setting the `init.ora` parameter `SHARED_POOL_SIZE` to 64 MB.

Answer: A

Explanation:

To decrease the response time to 3 seconds you need to be sure that the buffer cache hit ratio is greater than 95 percent. Otherwise decrease the `SHARED_POOL_SIZE` to 32 MB and allocate the remaining 16 MB to the buffer cache by setting the `DB_CACHE_SIZE` parameter

Incorrect Answers

B: This change will not decrease of OLTP application time to respond to 3 seconds.

C: By increasing SHARED_POOL_SIZE to 64 MB you will not achieve 3-seconds respond time goal.

D: This change will not decrease of OLTP application time to respond to 3 seconds.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 275-277

Chapter 5: Tuning the Database Buffer Cache

QUESTION 190:

Which four statements about the database storage hierarchy are true? (Choose four)

A. One or more extents make up a segment.

B. In the Oracle database the block is the smallest unit of data file I/O.

C. In the Oracle database the block is the smallest unit of space that can be allocated.

D. In the Oracle database the granule is the smallest unit of space that can be allocated.

E. An extent is a logical unit of database storage space allocation made up of a number of contiguous data blocks.

F. A data block is set of extents that contains all the data for a specific logical storage structure within a table.

Answer: A, B, D, E

Explanation:

It is true that one or more extents make up a segment. In the Oracle database the block is the smallest unit of data file I/O that can be allocated. In the Oracle database the granule is the smallest unit of space that can be allocated. An extent is a logical unit of database storage space allocation made up of a number of contiguous data blocks.

Incorrect Answers

C: In the Oracle database the block is the smallest unit of data file I/O, the granule is the smallest unit of space that can be allocated.

F: An extent is a logical unit of database storage space allocation made up of a number of contiguous data blocks.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 445-446

Chapter 8: Tuning Disk I/O

QUESTION 191:

You have several tables that are updated frequently. These updates are populating null columns with data.

Which sizing would be least likely to cause migration?

A. DB_BLOCK_SIZE=8192;PCTFREE5

B. DB_BLOCK_SIZE=8192;PCTFREE5

C. DB_BLOCK_SIZE=8192;PCTFREE10

D. DB_BLOCK_SIZE=8192;PCTFREE0

Answer: B

Explanation:

Setting PCTFREE to 20 least likely will cause migration. You need to set this value relatively high to preserve some space for future updates.

Incorrect Answers

A: This PCTFREE value will not be enough to prevent a row migration.

C: PCTFREE = 10 would be enough in some cases, but it is not the best DB_BLOCK_SIZE, PCTFREE configuration.

D: Setting PCTFREE to 0 will cause extensive row migration due to null columns updates.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 445-446
Chapter 8: Tuning Disk I/O

QUESTION 192:

To minimize fragmentation of your shared pool space, when is the best time for you to pin objects with the DBMS_SHARED_POOL package?

A. When the object is used for the first time.

B. At instance startup; load the object into memory and pin it

C. Immediately after instance startup, before loading anything into memory.

D. When the V\$LIBRARYCACHE view contains higher values in the RELOADS column than in the PINS column.

E. When the sum of values in the SHARABLE_MEMORY column of the V\$DB_OBJECT_CACHE view exceeds the value of the SHARED_POOL_SIZE initialization parameter.

Answer: B

Explanation:

These large and often-used objects (blocks of code in shared pool) should be pinned in shared pool immediately after startup, so as to reduce fragmentation. It is important to remember that an object must be read into the shared pool before it can be pinned.

Incorrect Answers

A: An object must be read into the shared pool at instance startup, not when it is used for the first time.

C: It's not possible to pin object after instance startup before loading anything into memory: at least data dictionary objects are already loaded into memory.

D: It's not related directly with V\$LIBRARYCACHE view values.

E: It's not related directly with V\$DB_OBJECT_CACHE view or SHARED_POOL_SIZE initialization parameter values.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 210-213
Chapter 4: Tuning the Shared Pool

QUESTION 193:

Which commands force DBWR to write dirty buffers to disk? (Choose three)

- A. COMMIT;
- B. ROLLBACK;
- C. SHUTEDDOWN IMMEDIATE;
- D. ALTER TABLESPACE A BEGIN BACKUP;
- E. ALTER TABLESPACE B OFFLINE NORMAL;

Answer: C, D, E

Explanation:

By shutting the database down you force DBWR to write dirty buffers to disk. Also this writing procedure is happening if you put a tablespace into BACKUP or OFFLINE mode.

Incorrect Answers

A: There is no CMOOT command in Oracle.

B: This command is used to rollback DML operations, it does not force DBWR to write dirty buffers .

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 250-252
Chapter 5: Tuning the Database Buffer Cache

QUESTION 194:

Which option accurately describes the buffer cache pools? (Choose three)

- A. DEFAULT
- B. INIT
- C. KEEP
- D. RECYCLE
- E. RETAIN
- F. PINNED

Answer: A, C, D

Explanation:

Unless you specify otherwise in the init.ora, only the Default Pool is created at instance startup. But Oracle provides you also with the ability to divide the Database Buffer Cache into as many as three separate areas called Buffer Pools: Keep Pool (optional), Recycle Pool (optional) and Default Pool (mandatory).

Incorrect Answers

B: There is no INIT buffer cache pool in the Database Buffer Cache.

E: There is no RETAIN buffer cache pool in the Database Buffer Cache.

F: There is no PINNED buffer cache pool in the Database Buffer Cache.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 265-266,

468-476

Chapter 5: Tuning the Database Buffer Cache, Chapter 9: Tuning Contention

QUESTION 195:

How to set the OPTIMIZER_MODE for minimizing total response time?

- A. RULE
- B. ALL_ROWS
- C. FIRST_ROWS_n
- D. FIRST_ROWS

Answer: B

Explanation:

The optimizer uses a cost-based approach for all SQL statements in the session regardless of the presence of statistics and optimizes with a goal of best throughput (minimum resource use to complete the entire statement). By setting the OPTIMIZER_MODE into ALL_ROWS you will minimize TOTAL response time. The CBO is a statistics-driven optimizer, and mode ALL_ROWS, FIRST_ROWS, FIRST_ROWS_n, CHOOSE are CBO settings.

Incorrect Answers

A: The optimizer chooses a rule-based approach for all SQL statements regardless of the presence of statistics. It does not minimize total response time.

C: The optimizer uses a cost-based approach, regardless of the presence of statistics, and optimizes with a goal of best response time return the first n number of rows; n can equal 1, 10, 100, or 1000.

D: The optimizer uses a mix of cost and heuristics to find a best plan for fast delivery of the first few rows.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 99-101

Chapter 3: SQL Application Tuning and Design

QUESTION 196:

Which statements are true about the Oracle block size? (Choose four)

- A. Large block size is good for DSS.
- B. Large block size may increase index reads.
- C. Small block size may decrease index reads.
- D. Large block size decreases contention.
- E. Microsoft SQL Server is better than Oracle by far.
- F. Large block size may waste space in the buffer cache.

Answer: A, B, C, F

Explanation:

DSS database will work better with a large database block size due to lot of full scans. Users of these systems are concerned with response time, which is the time it takes to get the results from their queries, so data need to be packed as closely as possible into blocks. Random access (index reads) to large object does not prefer a large block size: only consecutive access can benefit from a large block. And conversely: small block size may decrease index reads. Also it's true that large block size may waste space in the buffer cache.

Incorrect Answers

D: Large block size increases, not decreases, contention, since more data is locked when using larger block size. So large blocks are good for DSS systems and small ones - for OLTP where contention is high.

E: By choosing suitable block size you can achieve more effective performance results than Microsoft SQL Server.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 187-188
Chapter 3: SQL Application Tuning and Design

QUESTION 197:

Application Development and database configuration are closely related when a DBA needs to tune the database.

Which statement correctly reflects a dependency that exists in these two areas?

- A. Highly normalized tables will provide optimal query performance.
- B. Uncommitted SQL are unlikely to effect database performance.
- C. Every type of SQL has the same impact on DB Performance.
- D. Tuning index may help avoid excessive number of Full Table Scans.
- E. The impact of adding a new application to the system is directly measurable by doubling the workload.

Answer: D

Explanation:

Your main goal during tuning is to avoid a FTS, full-table scans, for large tables. By using indexes you can achieve this goal.

Incorrect Answers

A: Highly normalized tables will not always provide optimal query performance. Sometimes you need to decrease normalization level to achieve better performance results.

B: Uncommitted SQL statements are effect database performance because you need to maintain rollback segments for a read-consistent copy of data.

C: Every type of SQL has different impact on DB performance. For example, when SELECT statement uses FTS it can be disastrous for large tables, but it will be very good for small lookup tables. And vice versa, each DML operation will be slower for table with indexes than for flat table.

E: The impact of adding a new application to the system cannot be directly measurable by doubling the workload. It depends on a lot of other factors.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 187-188
Chapter 3: SQL Application Tuning and Design

QUESTION 198:

You are tuning SQL inside an application, which utility will collect detail info regarding the execution of SQL statements?

- A. CKPROF
- B. UTLBSTAT/UTLESTAT
- C. SQL TRACE
- D. Log Miner
- E. EXPLAIN PLAN

Answer: C

Explanation:

SQL TRACE is used to collect detailed info regarding the execution of SQL statements.

Incorrect Answers

A: CKPROF is used to format the trace file generated by SQL_TRACE.

B: UTLBSTAT /UTLESTAT utilities are used to gather database performance snapshots. They are depreciated after introduction of STATSPACK package for performance analysis.

D: Log Miner utility can be helpful to analyze online and archived redo logs.

E: EXPLAIN PLAN shows the most possible execution plan for SQL statement based on the optimizer mode, which, however, can differ from real one used for actual SQL statement execution. SQL TRACE provides you with plan really used during execution.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 40-41

Chapter 2: Sources of Tuning Information

QUESTION 199:

What should you set a small value in a Shared Server environment, because Oracle automatically adapts it?

- A. PROCESSES
- B. MAX_SHARED_SERVERS
- C. SHARED_SERVERS

Answer: C

Explanation:

SHARED_SERVERS is the total number of shared servers started when the database is started. It also represents the total number of shared servers PMON tries to keep. It should be the total number of servers expected to be used when the database is active.

Incorrect Answers

A: PROCESSES init.ora parameter limits maximum number of open sessions to the database.

B: MAX_SHARED_SERVERS parameter sets the upper bound on the total number of shared servers PMON can create based on the peak load of the system. This should be set high enough so that all requests can be serviced, but not so high that the system swaps if they are reached. The purpose of this parameter is to prevent the server from swapping. MAX_SHARED_SERVERS is intended to handle peak load.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 517-547
Chapter 10: Operating System Tuning

QUESTION 200:

You execute this query:

```
SELECT DECODE(SUM(totalq),0,'No Responses',  
SUM(wait)/SUM(totalq)) "Average wait time"  
From v$queue q,v$dispatcher d  
WHERE q.type = 'DISPATCHER'  
AND q.paddr = d.paddr;
```

Which statement is true?

- A. A steadily increasing value indicates a possible need for more dispatchers.
- B. A steadily increasing value indicates a possible need for fewer dispatchers.
- C. This query does not indicate whether user processes are waiting for dispatchers.
- D. The query returns the percentage of time the dispatcher process of each protocol are busy.

Answer: A

Explanation:

If average wait time is steadily increasing this indicates most likely a need for more dispatchers.

Incorrect Answers

B: By decreasing a number of dispatchers you increase average waiting time.

C: This query indicates that user processes are waiting for dispatchers.

D: This query indicates that user processes are waiting for dispatchers.

Oracle OCP Oracle 9i Database:

Performance Tuning Exam Guide - Charles

A. Pack - Oracle Press

Page 262 Tuning Oracle Shared Server

QUESTION 201:

Your Oracle9i database is running on a server that has 4 CPUs, and has these settings for initialization parameters:

db_cache_size= 4M

shared_pool_size= 52428800

db_block_size= 2048
large_pool_size= 614400
java_pool_size = 20971520
db_file_multiblock_read_count= 8
parallel_max_servers= 5

What is a valid statement about the above scenario?

- A. Tune the number of latches by setting it to db_cache_size(50*db_block_size).
- B. Tune the number of latches by setting to large_pool_size/1000.
- C. Tune the number of latches by setting it to Number of CPUs *2 *3.
- D. Oracle9i would automatically calculate the number of latches required, based on the environment defined by the initialization parameters for the database and the operating system level parameters.

Answer: D

Explanation:

Oracle9i automatically calculates the number of latches.

Incorrect Answers

A: You don't need to tune the number of latches.

B: This query indicates that user processes are waiting for dispatchers.

C: You don't need to tune the number of latches.

Oracle OCP Oracle 9i Database:

Performance Tuning Exam Guide - Charles

A. Pack - Oracle Press

Page 28 Oracle Performance Tuning Methodology

QUESTION 202:

You execute the following query to diagnose the performance of each shared server.

```
SQL > SELECT name "NAME", paddr, requests,  
2 > (busy / (busy+idle)) * 100 "%TIME BUSY", status  
3> FROM VSSHARED_SERVER;
```

Based on the above query, you obtain the following result:

```
NAME PADDR REQUESST %TIME BUSY STATUS  
-----  
S000 0107D73B 51525 9.19084132 WAIT (RECEIVE)  
S001 0107B233 26817 5.07654792 WAIT (COMMON)  
S002 0107B3BE 6362 1.44008509 WAIT (RECEIVE)  
S006 0108574C 54 86.9953920 WAIT (RECEIVE)  
S008 0107B549 1 99.9994096 WAIT (ENQ)
```

What can you infer from this result?

- A. Shared server S001 must be handling a batch process.
- B. The user assigned to shared server S008 should probably have a dedicated server.
- C. Shared server S002 should probably be servicing a DSS type of application.

- D. Shared server S006 must be representing most of the client sessions.
- E. A DBA is probably assigned to shared server S000 because it is handling many requests.

Answer: B

Explanation:

Since shared server S008 shows 99.9994096 percent of time busy with status 'ENQ', the user assigned to this shared server should have a dedicated server that will be able to process user's query more efficiently than shared server.

Incorrect Answers

A: It looks like shared server S008, not S001, must be handling a long-running batch process.

C: Shared server S008 should probably be servicing a DSS type of application.

D: Shared server S006 is not representing most of the client sessions: it has 86.99% of busy time, but S008 has more busy time.

E: S000 is not handling many requests.

Oracle OCP Oracle 9i Database:

Performance Tuning Exam Guide - Charles

A. Pack - Oracle Press

Page 263 Tuning Oracle Shared Server

QUESTION 203:

A developer at Certkiller has sent you multiple user SQL trace files related to the user session.

Which Oracle utility will convert this trace data into useable information?

- A. CKPROF
- B. UTLESTAT
- C. UTLBSTAT
- D. STATSPACK
- E. LOG MINER
- F. SQL ANALYZER

Answer: A

Explanation:

CKPROF is used to format the trace file generated by SQL_TRACE.

You can run the CKPROF program to format the contents of the trace file and place the output into a readable output file. Optionally, CKPROF can also:

1. Determine the execution plans of SQL statements.
2. Create a SQL script that stores the statistics in the database.

CKPROF reports each statement executed with the resources it has consumed, the number of times it was called, and the number of rows which it processed. This information lets you easily locate those statements that are using the greatest resource.

With experience or with baselines available, you can assess whether the resources used are reasonable given the work done.

Incorrect Answers

B: UTLESTAT is used to start statistics calculation: utility is obsolete after STATSPACK appearance.

C: UTLESTAT is used to finish statistics calculation: utility is obsolete after STATSPACK appearance.

D: STATSPACK is used to calculate statistics and generate statistics reports.

E: LOG MINER does not format trace files.

F: SQL ANALYZER is utility to optimize SQL queries.

Oracle OCP Oracle 9i Database:

Performance Tuning Exam Guide - Charles

A. Pack - Oracle Press

Page 56 Diagnostic and Tuning Tools

QUESTION 204:

You have been seeing poor performance for inserts into a new table.

You queried V\$WAITSTAT and V\$SYSTEM_EVENT and determined that there is free list contention in your database.

Next, you identified the segments by joining the DBA_SEGMENTS table and V\$SESSION_WAIT view to identify the file and block where waits are occurring.

What is the next step you should take?

A. If the file and block identify a data block within a table, consider increasing the number of free lists on the table using the ALTER TABLE command and specifying the FREELISTS clause.

B. If the file and block identify a data block within a table, consider increasing the number of free lists on the table using the ALTER TABLESPACE command and specifying the FREELISTS clause in the default storage clause.

C. If the file and block identify a segment header for a table, consider increasing the number of free lists on the table using the ALTER TABLE command and specifying the FREELISTS clause in the storage clause.

D. If the file and block identify a segment header for a table, consider increasing the number of free lists on the table using the ALTER TABLESPACE command and specifying the FREELIST clause in the default storage clause.

Answer: C

Explanation:

The number of free lists on the table should be increased if the file and block identify a segment header for a table using the ALTER TABLE command and specifying the FREELISTS clause in the storage clause.

Incorrect Answers

A: The number of free lists on the table should be increased if the file and block identify a segment header, not a data block within a table.

B: ALTER TABLE, not ALTER TABLESPACE, command should be used.

D: ALTER TABLE, not ALTER TABLESPACE, command should be used.

Oracle OCP Oracle 9i Database:

Performance Tuning Exam Guide - Charles

A. Pack - Oracle Press

Page 111 Sizing the buffer cache

QUESTION 205:

You want to enhance performance of the Database Optimizer to minimize total response time, thereby increasing overall throughput of batch process.

Which choice identifies the best optimizer mode setting for meeting this requirement?

A. RULE

B. ALL_ROWS

C. FIRST_ROWS

D. FIRST_ROWS_n

Answer: B

Explanation:

The optimizer uses a cost-based approach for all SQL statements in the session regardless of the presence of statistics and optimizes with a goal of best throughput (minimum resource use to complete the entire statement). By setting the OPTIMIZER_MODE into ALL_ROWS you will minimize TOTAL response time. The CBO is a statistics-driven optimizer, and mode ALL_ROWS, FIRST_ROWS, FIRST_ROWS_n, CHOOSE are CBO settings.

Incorrect Answers

A: The rule-based optimizer mode does not use table or index statistics at all. When RULE is specified, statistics are not used. In this case total response time will not be minimal.

C: By setting the OPTIMIZER_MODE into FIRST_ROWS you will minimize response time for first row.

D: By setting the OPTIMIZER_MODE into FIRST_ROWS_n you will minimize response time for first N rows.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 99-101
Chapter 3: SQL Application Tuning and Design

QUESTION 206:

Which dynamic view is most useful for determining the hit ratio for multiple buffer pools?

A. V\$CACHE

B. V\$SESS_IO

- C. V\$SYSSTAT
- D. V\$BUFFER_POOL
- E. V\$SYSTEM_EVENT
- F. V\$BUFFER_POOL_STATISTICS

Answer: F

Explanation:

V\$BUFFER_POOL_STATISTICS displays information about all buffer pools available for the instance. The "sets" pertain to the number of LRU latch sets.

Incorrect Answers

A: This is an Oracle9i Real Application Clusters view. This view contains information from the block header of each block in the SGA of the current instance as related to particular database objects.

B: This view lists I/O statistics for each user session.

C: The V\$SYSSTAT dynamic performance view contains statistics regarding overall system performance, gathering since instance startup.

D: This view also displays information about all buffer pools available for the instance. The "sets" pertain to the number of LRU latch sets. But it's not so useful as V\$BUFFER_POOL_STATISTICS.

E: This view contains information on total waits for an event. Note that the TIME_WAITED and AVERAGE_WAIT columns will contain a value of zero on those platforms that do not support a fast timing mechanism. If you are running on one of these platforms and you want this column to reflect true wait times, you must set TIMED_STATISTICS to TRUE in the parameter file ;doing this will have a small negative effect on system performance.

Oracle OCP Oracle 9i Database:

Performance Tuning Exam Guide - Charles

A. Pack - Oracle Press

Page 108 Sizing the buffer cache

QUESTION 207:

An application that uses your database keeps cursors open.

What should you confirm before changing the CURSOR_SPACE_FOR_TIME parameter in your database initialization file to TRUE?

- A. The TIMED_STATISTICS parameter is set to TRUE.
- B. The hit percentage in the buffer cache is at least 95%.
- C. The OPEN_CURSOR parameter is set to at least twice the default value.
- D. The value in the RELOADS column of V\$LIBRARYCACHE is consistently zero or close to zero.

Answer: D

Explanation:

When `CURSOR_SPACE_FOR_TIME` init.ora parameter is set to `TRUE`, shared SQL areas are pinned in the Shared Pool. This prevents the LRU mechanism from removing a shared SQL area from memory unless all cursors that reference that shared SQL area are closed. This can improve the Library Cache hit ratio and reduce SQL execution time, at the expense of using more of the server's main memory. The default value for `CURSOR_SPACE_FOR_TIME` is `FALSE`. To switch it to `TRUE` you should confirm that the value in the `RELOADS` column of `V$LIBRARYCACHE` is consistently zero.

Incorrect Answers

A: You don't need to check that the `TIMED_STATISTICS` parameter is set to `TRUE`. `TIMED_STATISTICS` specifies whether or not statistics related to time are collected.

B: It is not required to confirm that the hit percentage in the buffer cache is at least 95%.

C: The `OPEN_CURSOR` parameter setting to at least twice the default value is not necessary.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 232

Chapter 4: Tuning the Shared Pool

QUESTION 208:

In prior version of Oracle, when updating a child table `emp` (employees) the parent table `dept` (departments) implicitly received a Share Row Exclusive lock due to referential integrity.

What must you do when using Oracle9i?

A. You must turn on query rewrite for the instance.

B. You must index the foreign key column in the `emp` table.

C. You must place a unique index on the foreign key column in the `emp` table.

D. You must use a bitmapped index on the foreign key column in the `emp` table.

E. You do not need to do anything, the implementation of the foreign key constraint is modified.

Answer: E

Explanation:

You should not do anything because this restriction is moved away in Oracle9i.

Incorrect Answers

A: You don't need to turn on query rewrite for the instance.

B: It is not required to index the foreign key column in the `emp` table.

C: It is not required to index the foreign key column in the `emp` table, even a unique one.

D: It is not required to index the foreign key column in the `emp` table, even a bitmapped one.

Oracle OCP Oracle 9i Database:

Performance Tuning Exam Guide - Charles

A. Pack - Oracle Press

Page 239 Monitoring and detecting lock contention

QUESTION 209:

You are tuning the redo log buffer for an Oracle9i database.
Which four situations can affect the performance? (Choose four)

- A. DB_BLOCK_CHECKSUM is set to true.
- B. RMAN slave processes are too slow.
- C. Checkpoints are occurring every 2 minutes.
- D. Your junior DBA is creating large indexes in LOGGING mode.
- E. The ARCH process cannot complete the archive process quickly enough.

Answer: B, C, D, E

Explanation:

All these four situations can affect the performance.

Incorrect Answers

A: DB_BLOCK_CHECKSUM determines whether DBWn and the direct loader will calculate a checksum (a number calculated from all the bytes stored in the block) and store it in the cache header of every data block when writing it to disk.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 338-342
Chapter 7: Tuning Redo Mechanisms

QUESTION 210:

You can monitor and prevent locking problems by regularly monitoring which four views? (Choose four)

- A. V\$LOCK
- B. V\$SESSION
- C. V\$SYSSTATS
- D. DBA_WAITERS
- E. DBA_BLOCKERS
- F. V\$LOCKED_OBJECT

Answer: A, D, E, F

Explanation:

V\$LOCK, DBA_WAITERS, DBA_BLOCKERS and V\$LOCKED_OBJECT views can help you to identify and prevent locking problems.

Incorrect Answers

B: This view contains information about current sessions.

C: The V\$SYSSTAT dynamic performance view contains statistics regarding overall system performance, gathering since instance startup

Oracle OCP Oracle 9i Database:

Performance Tuning Exam Guide - Charles

A. Pack - Oracle Press
Page 241, 242 Monitoring and detecting lock contention

QUESTION 211:

Which are three possible causes of lock contention? (Choose three)

- A. Developers coding unnecessarily long transactions.
- B. Developers coding unnecessarily high locking levels.
- C. Developers using non-unique indexes on deferrable constraints.
- D. Applications converted from a non-Oracle database imposing higher locking levels.

Answer: A, B, D

Explanation:

If developers coding unnecessarily long transactions or use unnecessarily high locking levels or an applications converted from a non-Oracle database imposing higher locking levels it can cause a lock contention.

Incorrect Answers

C: Usage of non-unique indexes on deferrable constraints can decrease, not increase, a lock contention.

Oracle OCP Oracle 9i Database:

Performance Tuning Exam Guide - Charles

A. Pack - Oracle Press

Page 240 Monitoring and detecting lock contention

QUESTION 212:

You determined that the value for REQUEST_FAILURES as seen from V\$SHARED_POOL_RESERVED is more than zero and always increasing. Which action would be appropriate?

- A. Increase the value for the LARGE_POOL_SIZE parameter.
- B. Decrease the value for the LARGE_POOL_SIZE parameter.
- C. Increase the value for the SHARED_POOL_RESERVED_SIZE parameter.
- D. Decrease the value for the SHARED_POOL_RESERVED_SIZE parameter.

Answer: C

Explanation:

If REQUEST_FAILURES value is more than zero and increasing you have to increase the value for the SHARED_POOL_RESERVED_SIZE parameter.

SHARED_POOL_RESERVED_SIZE specifies (in bytes) the shared pool space that is reserved for large contiguous requests for shared pool memory. You can use this parameter to avoid performance degradation in the shared pool in situations where pool fragmentation forces Oracle to search for and free chunks of unused pool to satisfy the

current request.

Incorrect Answers

A: `LARGE_POOL_SIZE` lets you specify the size (in bytes) of the large pool allocation heap. The large pool allocation heap is used in shared server systems for session memory, by parallel execution for message buffers, and by backup processes for disk I/O buffers. (Parallel execution allocates buffers out of the large pool only when `PARALLEL_AUTOMATIC_TUNING` is set to true.)

B: You should consider `SHARED_POOL_RESERVED_SIZE`, not `LARGE_POOL_SIZE` parameter.

D: By decreasing `SHARED_POOL_RESERVED_SIZE` parameter performance will be forced to degrade.

Oracle OCP Oracle 9i Database:

Performance Tuning Exam Guide - Charles

A. Pack - Oracle Press

Page 75 Sizing the Shared Pool

QUESTION 213:

Which four are likely to be the cause of poor database performance? (Choose four)

- A. High occurrence of disk sorting.
- B. Low occurrence of recursive SQL.
- C. Too few redo logs that are too small.
- D. High use of cursors with bind variables.
- E. Limited use of cursors with bind variables.
- F. Frequent connections/disconnections from the application to the database.

Answer: A, C, E, F

Explanation:

High occurrence of disk sorting, few small redo logs, limited use of cursors with bind variables and frequent connection/disconnection from the application can cause of poor database performance.

Incorrect Answers

B: Low occurrence of recursive SQL will increase performance.

D: High use of cursors with bind variables in most cases increase performance due to less serialization.

Oracle OCP Oracle 9i Database:

Performance Tuning Exam Guide - Charles

A. Pack - Oracle Press

Page 28 Oracle Performance Tuning Methodology

QUESTION 214:

You are importing a table and you want to ensure that the table is put into a single contiguous area. There is not a single contiguous area large enough for the entire

1Z0-033

space allocated to the table.

Which two actions below could help you achieve your goal? (Choose two)

- A. Use the cascade option.
- B. Use the constraint option.
- C. Use the primary key option.
- D. Coalesce the tablespace the table is using.
- E. Add a datafile large enough to meet the allocation space required to the tablespace the table is using.

Answer: D, E

Explanation:

To avoid data migration during import you need coalesce free space or add a new datafile to tablespace. The RECORDLENGTH parameter is necessary when you must transfer the export file to another operating system that uses a different default value. The COMPRESS option specifies how Export and Import manage the initial extent for table data.

Incorrect Answers

A: There is no CASCADE option for export/import utilities.

B: Constraint option is related table constrains: it has nothing to do with space for table.

C: There is no PRIMARY KEY option for export/import utilities.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 215:

Which statement demonstrates a correct understanding of common tuning problems?

- A. It is generally recommended to host OLTP and DSS applications on a single database instance.
- B. Once your performance measurement criteria are defined and established in the development environment, they should directly apply to the production environment.
- C. Generic data modeling and highly normalized schema implementation will generally provide long-term system stability and optimal performance.
- D. If your application is known to be CPU-bound and the database has exhibited serious system resource starvation, it is generally worthwhile to consider moving procedural business logic from the database layer to another layer, such as middle-tier.
- E. With the rapid advance of hardware technology, performance issues are generally handles at the hardware level and no benefits would be realized by tuning software components such as the database and applications.

Answer: D

Explanation:

Usually it's good idea to implement multi-tier application, including application server, if your application is CPU-bound and the database has exhibited serious system resource starvation .

Incorrect Answers

A: It is not recommended to host OLTP and DSS applications on a single database instance.

B: You should test application in QA (and/or Staging) environment before moving it into Production.

C: Generic data modelling and high normalized schema implementation usually do not provide good performance: it's a trade of de-normalization and performance. In reality nobody uses 5th normal form.

E: The best result in tuning you can achieve by tuning software components such as the database and applications. Advance on hardware level will not give you strong improvement and it usually does not fix the core of performance problem in long-term run.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 216:

As a conscientious Certkiller DBA you are aware that tuning should begin at the initial stage of the application development phase.

Assuming you have optimal data model and data file layout to eliminate hot spots, what is the next step you should take to isolate the problems that may cause poor database performance?

- A. Tune indexes
- B. Increase SGA
- C. Add more CPUs
- D. Add more LRU latches
- E. Monitor performance

Answer: E

Explanation:

Before any tuning efforts you need to monitor performance for a while, noticing database behavior in high-pick and normal hours. You should also to find your "low hanging fruit" (usually it's SQL query consuming the most of the database resources).

Incorrect Answers

A: You should tune indexes if you see a lot of full table scans (FTS). However, in DSS environments FTS are usual and indexes in some cases will decrease performance.

B: It's not your first step in tuning.

C: Hardware tuning is usually the last step: the biggest gain you can achieve by application (SQL) tuning.

D: LRU latches should be tuned when you experience high level of serialization and only in case you already tried SQL tuning to decrease latch usage.

QUESTION 217:

There are some parameters DBAs need to set in init<SID>.ora while configuring the Oracle Shared Server.

Which parameter should be set to a small value because Oracle dynamically adjust the value based on the length of the request queue?

- A. DISPATCHERS
- B. SHARED_SERVERS
- C. MAX_DISPATCHERS
- D. LARGE_POOL_SIZE
- E. MAX_SHARED_SERVERS
- F. SHARED_SERVER_SESSIONS

Answer: B

Explanation:

SHARED_SERVERS is the total number of shared servers started when the database is started. It also represents the total number of shared servers PMON tries to keep. It should be the total number of servers expected to be used when the database is active.

Incorrect Answers

A: DISPATCHERS configures dispatcher processes in the shared server architecture. Oracle does not dynamically adjust its value based on the length of the request queue.

C: MAX_DISPATCHERS specifies the maximum number of dispatcher processes allowed to be running simultaneously. The default value applies only if dispatchers have been configured for the system. The value of MAX_DISPATCHERS should at least equal the maximum number of concurrent sessions divided by the number of connections for each dispatcher.

D: Oracle does not dynamically adjust LARGE_POOL_SIZE value based on the length of the request queue.

E: MAX_SHARED_SERVERS specifies the maximum number of shared server processes allowed to be running simultaneously. If artificial deadlocks occur too frequently on your system, you should increase the value of MAX_SHARED_SERVERS.

F: SHARED_SERVER_SESSIONS specifies the total number of shared server architecture user sessions to allow. Setting this parameter enables you to reserve user sessions for dedicated servers.

Oracle OCP Oracle 9i Database:

Performance Tuning Exam Guide - Charles

A. Pack - Oracle Press

Page 263 Tuning Oracle Shared Server

QUESTION 218:

What is the initialization parameter PGA_AGGREGATE_TARGET used for?

- A. Specifying the number of cached tables to preload at instance startup.
- B. Specifying the maximum number of parallel execution processes and parallel recovery processes for an instance.
- C. Specifying the maximum number of operating user processes that can simultaneously connect to Oracle.
- D. Enabling the automatic sizing of SQL working areas used by memory-intensive SQL operators such as sort, group-by, hash-join, bitmap merge and bitmap create.

Answer: D

Explanation:

PGA_AGGREGATE_TARGET specifies the target aggregate PGA memory available to all server processes attached to the instance. You must set this parameter to enable the automatic sizing of SQL working areas used by memory-intensive SQL operators such as sort, group-by, hash-join, bitmap merge, and bitmap create.

Incorrect Answers

- A: It does not specify the number of cached tables to preload at instance startup.
- B: It does not specify the maximum number of parallel execution processes and parallel recovery processes for an instance.
- C: It does not specify the maximum number of operating user processes that can simultaneously connect to Oracle.

Oracle OCP Oracle 9i Database:

Performance Tuning Exam Guide - Charles

A. Pack - Oracle Press

Page 170 Optimize Sort Operations

QUESTION 219:

You work as a programmer at Certkiller Inc. During the application development phase, you receive complaints from the Certkiller application developers that the database is slow while queuing certain tables. The application is a DSS application that has been in development for the past 6 months, but the database slowdown has been experienced only recently.

The server is dedicated to one Oracle database instance that hosts the application, and has 1 gigabyte of physical memory. You followed memory-tuning guidelines and allocated approximately 164 megabytes for the total SG

- A. Using one of the database performance diagnostic tools, you confirmed the following:
1. The statistics are captured for the duration of 1 hour.
 2. Buffer cache hit ratio is average of 73%.
 3. Library cache hit ratio is average of 98%.
 4. Disk sort ratio is average of 0.2% per transaction.

1Z0-033

5. Sum of the long table scans is average of 13.5 per transaction.
6. Sum of the physical block reads of each tablespace is average of 0 with the exception of one tablespace that has the value of 2306797.

What is the first action you should consider?

- A. Add more LRU latches in order to reduce process contention.
- B. Add more disks to distribute data files in order to eliminate file I/O contention.
- C. Review the application code and look for the cause of high physical block read count centered on one particular tablespace.
- D. Look for literal SQLs in the application code and replace them with bind variables wherever possible.
- E. Increase SORT_AREA_SIZE in the parameter file to provide more memory space for sort operations such as ORDERED BY, GROUP BY, and DISTINCT.

Answer: C

Explanation:

You need to find your "low hanging fruit", SQL query consuming the most of the database resources. In your case it's query causing of high physical block read count centered on one particular tablespace.

Incorrect Answers

A: LRU latches should be tuned when you experience high level of serialization and only in case you already tried SQL and memory tuning to decrease latch usage.

B: Hardware tuning is usually the last step: the biggest gain you can achieve by application (SQL) tuning.

D: You need to tune only one SQL query, consuming the most of the database resources, your "low hanging fruit", not entire application code. Also it's known that literal usage for DSS application can be more helpful than bind variables.

E: Memory should be tuned after SQL tuning.

Oracle OCP Oracle 9i Database:

Performance Tuning Exam Guide - Charles

A. Pack - Oracle Press

Page 344 SQL Statement Tuning

QUESTION 220:

Which statement describes how execution plans can be stabilized across Oracle releases, database changes, and other factors that normally cause an execution plan to change?

- A. You can create a stored outline containing a set of hints used by the optimizer to create an execution plan.
- B. You can create a stored explain plan containing a set of hints used by the optimizer to create an execution plan.
- C. Oracle9i provide you with a method to write the code to detect which optimizer is in use so the code can take advantage of the improved execution plan.
- D. Oracle9i automatically recompiles the code the first time the code is executed.

This ensures the code will always use the most current execution plan.

Answer: A

Explanation:

You can create a stored outline containing a set of hints used by the optimizer to create an execution plan. Outlines can be categorized to simplify the management task. The CREATE OUTLINE statement allows for specification of a category. The DEFAULT category is chosen if unspecified. Likewise, the CREATE_STORED_OUTLINES parameter lets you specify a category name, where specifying true produces outlines in the DEFAULT category.

Incorrect Answers

B: You can create a stored outline, not explain plan.

C: The code is optimizer-independent.

D: By recompiling the code you Oracle can choose different execution plan, which does not stabilize it across Oracle releases.

Oracle OCP Oracle 9i Database:

Performance Tuning Exam Guide - Charles

A. Pack - Oracle Press

Page 335 SQL Statement Tuning

QUESTION 221:

In order to avoid validation errors, every resource plan directive must include which resource consumer plan?

A. OTHER

B. SYS_GROUP

C. OTHER_GROUPS

D. DEFAULT_GROUPS

Answer: C

Explanation:

This description fits the OTHER_GROUPS consumer group.

DEFAULT_CONSUMER_GROUP is the initial consumer group for all users/sessions that have not been explicitly assigned an initial consumer group.

DEFAULT_CONSUMER_GROUP has switch privileges granted to PUBLIC therefore, all users are automatically granted switch privilege for this consumer group.

Incorrect Answers

A: There is OTHER_GROUPS, not OTHER, resource consumer plan in Oracle.

B: SYS_GROUP is the initial consumer group for the users SYS and SYSTEM.

LOW_GROUP provides a group having lower priority than SYS_GROUP and OTHER_GROUPS in this plan.

D: There is DEFAULT_CONSUMER_GROUP, not DEFAULT_GROUPS resource consumer plan in Oracle.

Oracle OCP Oracle 9i Database:
Performance Tuning Exam Guide - Charles
A. Pack - Oracle Press
Page 365 OS Considerations and Oracle Resource Manager

QUESTION 222:

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

Explanation:

It's recommended to separate STATSPACK data from Production data to avoid statistics gathering influence to production statistics. It can be used to compare performance from one day, week, or year to the next, there must be multiple snapshots taken over a period of time. 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: it allows you to drill down your "low hanging fruit" SQL queries.

Incorrect Answers

B: Obsolete statistical data collected by STATSPACK can be purged with SPPURGE.SQL script.

E: STATSPACK will not automatically generate a report: you need to run SPREPORT.SQL script to see statistical output for instance level or SPREPSQL.SQL for particular SQL query.

Oracle OCP Oracle 9i Database:
Performance Tuning Exam Guide - Charles
A. Pack - Oracle Press
Page 51 Diagnostic and tuning Tools

QUESTION 223:

Which three SQL statements will cause the DBWR process to write dirty blocks back to the data files? (Choose three)

- A. COMMIT;
- B. ROLLBACK;
- C. ALTER SYSTEM FLUSH SHARED_POOL;
- D. SHUTEDDOWN IMMEDIATE;
- E. ALTER TABLESPACE TSP_A OFFLINE
- F. ALTER TABLESPACE TSP_B BEGIN BACKUP;

Answer: D, E, F

Explanation:

These statements will force DBWR process to write dirty blocks back to the data files. By shutting the database down you force DBWR to write dirty buffers to disk. Also this writing procedure is happening if you put a tablespace into BACKUP or OFFLINE mode.

Incorrect Answers

A: COMMIT statement does not force DBWR to write dirty blocks: LGWR writes changes into online redo logs.

B: ROLLBACK statement does not force DBWR to write dirty blocks: LGWR writes changes into online redo logs.

C: By flushing shared pool it's not required for DBWR to write dirty blocks.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 390-391

Chapter 8: Tuning Disk I/O

QUESTION 224:

Deadlocks are occurring in the database.

Where can you find information about these deadlocks?

- A. Audit trail
- B. Online redo log
- C. LGWR trace files
- D. Alert log and foreground server trace files

Answer: D

Explanation:

Alert log and foreground server trace files contain information about these deadlocks.

Incorrect Answers

A: Audit trail is not used for this purpose.

B: Online redo log does not have any information about deadlocks.

C: LGWR trace files do not have any information about deadlocks.

Oracle OCP Oracle 9i Database:

Performance Tuning Exam Guide - Charles

A. Pack - Oracle Press

Page 245 - 247 Monitoring and Detecting lock Contention

QUESTION 225:

Which statement is true about the SGA in Oracle9i?

- A. The unit of allocation for resizing the buffer cache or shared pool is defined by the init.ora parameter DB_BLOCK_SIZE.
- B. The unit of allocation for resizing the buffer cache or shared pool within the SGA is defined by the size of the granules in use for the current SGA.
- C. The maximum size of the SGA is set by the init.ora parameter SGA_MAX_SIZE and can be altered dynamically.
- D. You can use a combination of the init.ora parameters SHARED_POOL_SIZE, DB_BLOCK_BUFFERS and SGA_MAX_SIZE to define the size of a dynamic SGA.

Answer: B

Explanation:

In the Oracle database the granule is the smallest unit of space that can be allocated. The unit of allocation for resizing the buffer cache or shared pool within the SGA is defined by the size of the granules in use for the current SGA.

Incorrect Answers

A: The unit of allocation for resizing the buffer cache or shared pool is defined by the size of the granules in use for the current SGA, not by the init.ora parameter DB_BLOCK_SIZE.

C: SGA_MAX_SIZE cannot be altered dynamically.

D: This statement is wrong.

Oracle OCP Oracle 9i Database:

Performance Tuning Exam Guide - Charles

A. Pack - Oracle Press

Page 101 Sizing the Buffer Cache

QUESTION 226:

Which four statements are true regarding the size and the effect of Oracle Blocks?
(Choose four)

- A. Large Oracle blocks are good for DSS types of application.
- B. Larger Oracle blocks may increase the performance of the index reads.
- C. Smaller Oracle blocks may decrease the performance of the index reads.
- D. Small Oracle blocks may increase block contention, because there are fewer rows per block.
- E. Larger Oracle blocks may waste the space in the buffer cache if the applications randomly access rows.

Answer: A, B, C, E

Explanation:

DSS database will work better with a large database block size due to lot of full scans. Users of these systems are concerned with response time, which is the time it takes to get the results from their queries, so data need to be packed as closely as possible into blocks. Random access (index reads) to large object does not prefer a large block size: only consecutive access can benefit from a large block. And conversely: small block size may decrease index reads. Also it's true that large block size may waste space in the buffer cache.

Incorrect Answers

D: Small Oracle blocks may decrease block contention, because there are fewer rows per block: there are few chances that some processes will try to read the same block simultaneously.

Oracle OCP Oracle 9i Database:

Performance Tuning Exam Guide - Charles

A. Pack - Oracle Press

Page 306 - 307 Using Oracle Block Efficiently

QUESTION 227:

Which two statements are true regarding running the Data Dictionary Cache?
(Choose two)

- A. The ratio of GETMISSES to GETS found in the V\$ROWCACHE view should be less than 5 percent.
- B. If the hit ratio for the library cache is acceptable, the hit ratio for the data dictionary should also be acceptable.
- C. On instance startup, the data dictionary cache contains no data, so any SQL statement is likely to result in cache misses.
- D. On instance startup the data dictionary cache is automatically loaded with the most commonly used information, so many SQL statements will not cause cache misses.

Answer: B, C

Answer A is not true because this percent usually should be < 15%

Answer D is not true - there is such definition "most commonly used information" while instance is starting

Reference: Oracle9i Performance Tuning - Student Guide, Lesson 3

QUESTION 228:

You have been running SQL that requires sorts all day in your database. At the end of the day, you queried the V\$SORT_SEGMENT view using this statement:

```
SELECT TABLESPACE_NAME, CURRENT_USERS, USED_EXTENDS,  
FREE_EXTENDS FROM V$SORT_SEGMENT
```

The statement returned NO rows.

What is a valid conclusion?

- A. All sort operations have completed in memory.

- B. All sort operations went to the temporary tablespace.
- C. The PRE_PAGE_SGA initialization parameter was set to TRUE.
- D. All the sort operations went to the users' default tablespace.

Answer: A

Reference: Oracle9i Performance Tuning - Student Guide, Lesson 7

QUESTION 229:

You are monitoring redo log performance using the query shown.

```
SQL > select sid, event seconds_in_wait, stat  
2 > from v$session_wait  
3> WHERE EVENT='LOG BUFFER SPACE';
```

This is returned:

```
SID EVENT SECONDS_IN_WAIT STATE  
-----
```

```
5 log buffer space 110 WAITING
```

Which two options could solve your problem? (Choose two)

- A. Make the log buffer bigger.
- B. Perform a commit more frequently.
- C. Increase the size of the log files.
- D. Add a disk controller to the file.
- E. Move the log files to a faster disk.

Answer: A, E

Explanation:

This wait event indicates that the lot of time is spent for waiting for enough space in the redo log buffer. Redo Log Buffers are filled faster then LGWR process is writing to online redo log files.

Reference: Oracle9i Performance Tuning - Student Guide, Lesson 5

QUESTION 230:

Which two are true about buffer cache advisory? (Choose two)

- A. It is enabled by the DB_CACHE_ADVICE initialization parameter.
- B. It estimates the number of physical writes to disk from the buffer cache.
- C. It predicts the estimated number of indirect reads for different cache sizes.
- D. It enables and disables statistics for predicting behavior with different cache sizes.
- E. The database should be shut down after enabling the DC_CACHE_ADVISOR to ensure that buffer cache statistics are reset.

Answer: A, D

Reference: Oracle9i Performance Tuning - Student Guide, Lesson 4

QUESTION 231:

Which three statements are true about the High-water mark? (Choose three)

- A. It is reset by the truncate command.
- B. It is recorded in segment header block.
- C. It is set to the end of the segment on creation.
- D. It is set to the beginning of the segment on creation.

Answer: A, B, D

Reference: Oracle9i Performance Tuning - Student Guide, Lesson 13

QUESTION 232:

Many applications insert in ascending index order and delete older values. In which case would you need to rebuild your indexes regularly?

- A. When the B-level is low.
- B. When index blocks contain many entries.
- C. When the high water mark has been reached.
- D. When the B-level is high and the index blocks contain very few entries.

Answer: D

Reference: Oracle9i Performance Tuning - Student Guide, Lesson 12

QUESTION 233:

Which of these are legitimate performance tuning goals? (Choose three.)

- A. Reduce throughput.
- B. Increase SQL run times.
- C. Lower hit ratios.
- D. Decrease sorts to disk.
- E. Eliminate lock contention.
- F. Optimize connection management.

Answer: D, E, F

We cannot consider reducing throughput as a correct choice because we want to increase throughput. We also want to decrease SQL run times and raise hit ratios. Sorts to disk are significantly slower than sorts in memory. We also want to reduce or eliminate lock contention so that user processes are not waiting. We also want to optimize connection management because poor connection management creates unnecessary overhead on the network and the database server.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 234:

The best place to start looking for general resource usage statistics is in which V\$ view?

- A. V\$INSTANCE
- B. V\$INSTANCE_STATS
- C. V\$SYSSTAT
- D. V\$PROCESS
- E. V\$WAITSTAT

Answer: C

V\$SYSSTAT is the view that shows general resource usage statistics for the instance. V\$INSTANCE simply shows the state of the current instance. V\$INSTANCE_STATS view does not exist. V\$PROCESS has row for each process in the instance, but does not indicate resource usage. V\$WAITSTAT is specific to buffer wait events, not general resource usage.

Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 235:

Which of these columns are used to calculate the library cache Hit Ratio?

- A. GETS and PINS
- B. PINS and PINHITS
- C. PINS and RELOADES
- D. GETMISSES and GETS
- E. GETS and RELOADS

Answer: C

The library cache hit ratio is calculated as the sum of PINS minus the sum of RELOADS, divided by the sum of PINS, and multiplied by 100.

Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 236:

You have used the ALTER command to set the DB_CACHE_ADVICE parameter to READY. What is now happening in the instance?

- A. Memory has been allocated in the shared pool for cache advice buffers.
- B. CPU utilization will significantly increase because Oracle is collecting more detailed statistics about the buffer cache.
- C. The buffer cache hit ratio will now increase.
- D. The V\$DB_CACHE_ADVICE view is now populated.

Answer: A

When DB_CACHE_ADVICE is set to ready, memory for the advice collection is configured in the shared pool. CPU utilization will not increase because no data gathering occurs until you set the value to ON. DB_CACHE_ADVICE makes no changes to the buffer cache size; it merely forecasts physical I/O change to the instance if the buffer cache is configured smaller or larger. V\$DB_CACHE_ADVICE view remains unpopulated until the advice parameter is set to ON.

Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 237:

You noticed that the redo buffer retries ratio is 5 percent. Which of the following is a viable remedy?

- A. Move the archive redo log file destination to a slower device.
- B. Decrease the size of the redo log buffer.
- C. Increase the size of the redo log buffer.
- D. Decrease the size of the redo log files.

Answer: C

When the number of redo buffer allocation retries is greater than 1 percent of the number of redo entries, then we should consider increasing the size of the redo log buffer. The redo buffer allocation retries are indicators of waits for space in the redo log buffer. Moving the archive redo log file destination to a slower device will not improve the Waits; it will increase waits for the redo log switch. If you decrease the size of the redo log buffer it will more than likely increase the number of retries. Decreasing the size of the redo log files will cause more frequent log switching and contribute to waits.

Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 238:

Which of the following is an indicator of inappropriate tablespace usage?

- A. HR objects are stored in the SYSTEM tablespace.
- B. UNDO segments and rollback segments are in separate tablespaces.
- C. User sort segments are created in a tablespace of type TEMPORARY.
- D. Rollback segments and user sort segments are in the same tablespace as user data.
- E. Indexes and tables are in the same tablespace.
- F. Stored procedures are in the SYSTEM tablespace.

Answer: A, D, E

Objects owned by users other than SYS should not be in the SYSTEM tablespace because they can cause fragmentation and contention. Rollback and sort segments should

not be in the same tablespace as user data; rollback and temporary segment should be in separate tablespaces. Indexes and tables should not be in the same tablespace. UNDO cannot share a tablespace with other type of segments, and rollback segments should be in their own tablespace. If user sort segments are created in a tablespace of type TEMPORARY it is an appropriate tablespace recommendation, and it is a fact that stored procedures are stored in the SYSTEM tablespace.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 239:

Which of the following queries will mostly likely require a sort-merge join?

- A. Two-table equijoin with no indexes
- B. Two tables equijoin with one index
- C. Two-table equijoin in which both tables are indexed on the join column
- D. SELECT DISTINCT on a single table

Answer: B

Because neither table has an index on the join column, Oracle must read and sort each table and then merge the two result sets together. Two-table equijoin with no indexes and two-table equijoin in which both tables are indexed on the join column will use a nested-loop join; SELECT DISTINCT will get a sort on the query, but no join will occur.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 240:

Select the statement that most accurately describes the function of a latch.

- A. A latch is responsible for protecting database tables from concurrent changes by multiple processes.
- B. A latch prevents two or more process from modifying the same data structure at the same time.
- C. The purpose of a latch is to prevent Oracle memory structures from growing too large for real memory.
- D. The purpose of a latch is to coordinate which background process will be used to read data blocks from data files.
- E. Latches are used to speed recovery following a database crash.

Answer: B

The function of a latch is to protect Oracle memory structures from being modified by more than one process at a time. A latch is not responsible for protecting database tables from concurrent changes by multiple processes. This statement more accurately describes a lock instead of a latch. To prevent Oracle memory structures from growing too large for real memory is not the responsibility of a latch, but it is the responsibility of the DBA.

Other choices are simply not functions associated with latches.
Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 241:

You need to switch from one tablespace (active) to another (target). Which of the following conditions would prevent this action?

- A. The target tablespace is smaller than the active tablespace.
- B. The target tablespace is defined as a temporary tablespace.
- C. There are transactions in the active tablespace.
- D. Queries need the undo images in the active tablespace for read consistency.
- E. The initialization parameter UNDO_MANAGEMENT is set to AUTO.

Answer: B

The target tablespace must be defined as an undo tablespace, not a temporary tablespace. The size of the target and active tablespaces doesn't matter. Active transactions in the active tablespace will not prevent the switch to a new tablespace. The active transactions will continue, but no new transactions will begin in the tablespace. Queries required the undo images in the active tablespace for read consistency do not prevent the switch from occurring. It is required for automatic undo management that the initialization parameter UNDO_MANAGEMENT should be set to AUTO, so it will not prevent the switch from occurring.

Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 242:

As the DBA, you'll need to execute which script to create the DBA_BLOCKERS view?

- A. Catproc.sql
- B. Utillockt.sql
- C. Utblock.sql
- D. Catlock.sql
- E. Catblock.sql
- F. Dba_blockers.sql

Answer: E

The Catblock.sql script should be run shortly after the database is created and after the Catproc.sql script is executed. Catblock.sql creates several views, one of which is the DBA_BLOCKERS view that the DBA can use to monitor locking issues. To run a Catproc.sql is a prerequisite, but it does not create the DBA_BLOCKERS view. Utillockt is a monitoring script that will not function until the Catblock.sql script has been executed once in the database. Utblock.sql, Catlock.sql and Dba_blockers.sql scripts do not exist.

QUESTION 243:

Which of these environments would benefit from moving from dedicated servers to shared servers?

- A. A Web-based environment, with potentially thousands of concurrent users
- B. A distributed OLTP system
- C. A traditional data warehouse
- D. A 25-user Unix database with local connections
- E. A batch processing system
- F. A DSS or reporting system

Answer: A, B

Explanation:

The shared servers environment is characterized by a large number of nonlocal users whose database usage profile could be described as few and far between. A data warehouse, a 25-user Unix database with local connections and a batch processing system have characteristics that would benefit from dedicated servers. A traditional data warehouse, the batch system, and the DSS system are all characterized by processes that query large amounts of data or insert large amounts of data. These data-intense systems perform better with dedicated servers. A 25-user Unix database with local connections answer is kinda tricky; Unix systems can benefit from shared servers, but a 25-user system with local connections should probably remain with dedicated servers.

Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 244:

In an OLTP environment, three tables share a key column value and are always joined together in queries. Which one of these storage structures is a good design choice?

- A. IOT
- B. Cluster
- C. Materialized view

Answer: B

The question describes the ideal situation for building a cluster. Since the tables are always joined together with the same key, this key would be an ideal cluster key. We now need to determine if a cluster or an index cluster is more appropriate. IOT is not the ideal storage structure for the situation described. Materialized view might be a good choice in a DSS environment because all the criteria are met except for the environment.

QUESTION 245:

Which of the following is true about allocating extents?

- A. Always let Oracle allocate extents dynamically.
- B. You should avoid dynamic extent allocation.
- C. You can monitor segments that are close to extending by querying DBA_EXTENTS.
- D. You should extend a table manually with the ALTER TABLE ... ALLOCATE NEXT command.
- E. You should extend a table manually with the ALTER TABLE ... AUTOEXTEND command.

Answer: B

Dynamic extent allocation can cause performance problems for application SQL statements that are inserting or updating. Also we don't want Oracle to dynamically extend segments. DBA_EXTENTS does not give us information about segments close to extending. The ALLOCATE NEXT and the AUTOEXTEND are not valid ALTER TABLE clauses.

QUESTION 246:

When creating a local statistics table with the DBMS_STAT.CREATE_STAT_TABLE procedure, which of the following fields are required? (Choice all that apply.)

- A. ownname
- B. stattab
- C. tblspace
- D. statown

Answer: A, B

The ownname and stattab fields are required. Ownname specifies in which schema to create the table specified with the stattab variable. Tblspace, is optional if you want to create the table in some tablespace other than the default for the user specified in ownname. Statown, is not a CREATE_STAT_TABLE parameter.

QUESTION 247:

From the following, choose the correct sequence of events to configure Resource Manager:

- A. Create a pending area, create a plan, create plan directives, and create consumer groups.
- B. Create a pending area, create consumer groups, create plan directives, and create a plan.
- C. Create a pending area, create a plan, create consumer groups, and create plan directives.

Answer: C

This specifies the correct sequence of events. We did not include validate and submit steps, not did we assign users or roles to consumer groups. However, the steps to create a pending area, create a plan, create consumer groups, and create plan directives are in correct order. Other answers are both incorrect because the sequence of events is incorrect.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 248:

The system alert log is written to which directory specified in the init.ora file?

- A. USER_DUMP_DEST
- B. BACKGROUND_DUMP_DEST
- C. CORE_DUMP_DEST
- D. ALERT_DUMP_DEST

Answer: B

Specified in the init.ora file, and also dynamically configurable, the BACKGROUND_DUMP_DEST specifies the location for dump files from server background processes, and also the alert log. The USER_DUMP_DEST is used for user process trace output, and the CORE_DUMP_DEST directory is used for core dumps. The ALERT_DUMP_DEST is not valid initialization parameter.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 249:

The STATSPACK report is similar in nature to which of the following?

- A. OEM Oracle Expert
- B. OEM Tuning Template
- C. SQL Trace and CKPROF
- D. UTLSTAT report.txt

Answer: D

The Statspack report is meant to replace UTLBSTAT/UTLESTAT and the resulting

report.txt. Statspack enables the retention of multiple snapshots of performance statistics in the database, so you can generate a Statspack report to calculate the difference between parameters taken at two snapshots.

Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 250:

The dictionary cache hit ratio measures the ratio of which two statistics from V\$ROWCACHE? (Choose two.)

- A. PINS
- B. RELOADS
- C. GETS
- D. GETMISSES
- E. KEEPS

Answer: C, D

The actual formula is $(\text{sum}(\text{gets}) - \text{sum}(\text{getmisses})) / \text{sum}(\text{gets})$. You should try to keep this ratio in the high 80's or better. The dictionary cache hit ratio is an indicator of how often the correct data dictionary information is found in memory when needed. GETS indicates the number of data dictionary requests; GETMISSES indicates the number of requests that were not met by reading the cache and required an I/O.

Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 251:

In the shared server environment, the UGA components user session data and cursor state information are stored in which memory structures?

- A. PGA
- B. VBA
- C. Shared pool
- D. Buffer cache
- E. Log buffer

Answer: C

In the shared server environment, some of the UGA components are moved into the shared pool. These components reside in the PGA when dedicated servers are used. The session stack space remains in the PGA regardless of server configuration.

Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 252:

1Z0-033

Dynamic SGA allocation enables the DBA to increase the size of the SGA components up to?

- A. MAX_SGA_SIZE
- B. SGA_MAX_SIZE
- C. SHARED_POOL_SIZE+BUFFER_CACHE_SIZE
- D. DB_BLOCK_SIZE

Answer: B

The SGA_MAX_SIZE parameter indicates the maximum size that you can dynamically increase the SGA components up to. The buffer cache, redo buffers, the Java pool, large pool, shared pool, and reserved pool all must fit inside the SGA_MAX_SIZE limit.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 253:

Freelist contention can be eliminated by using automatic segment-space management, which is designated when?

- A. At the table create time
- B. At the tablespace create time
- C. At the database create time
- D. At the query run time

Answer: B

Automatic segment-space management is designated at the time the tablespace is created by issuing the EXTENT MANAGEMENT LOCAL clause in the CREATE TABLESPACE command.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 254:

Which of these environmental influences would lead you to increase the size of the java pool? (Choose two.)

- A. Large client-side applications are written in Java.
- B. The application uses large Java stored procedures.
- C. Shared servers are used for the Java components.
- D. Dedicated servers are used.
- E. Application users are spread across multiple contents.

Answer: B, C

If the application uses large Java stored procedures, then a large JAVA_POOL setting will contribute to system performance. Also, if Java components are used in the database

and the database uses shared servers, some of the Java process memory structures are kept in the Java pool. All other answers do not directly affect the Java pool requirements.
Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 255:

Which of the following is not considered an important goal of the database configuration phase?

- A. Tuning disk I/O
- B. Tuning memory
- C. Tuning application code
- D. Tuning CPU utilization

Answer: C

At the database configuration phase of the application design process, we are mostly concerned with tuning memory, I/O, and CPU. The application code should have been already tuned in the development and test environments. There are usually user acceptance testing and post-implementation opportunities to tune the application SQL code.

Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 256:

You noticed that all SQL statements that have a join condition on one particular table appear to be running very slowly. Which of these steps should you take first to help identify the performance problem?

- A. Check to see if the database is running the CBO or RBO.
- B. Check file I/O statistics for the data file and tablespace that the table is in.
- C. Verify that there are no missing indexes from the table.
- D. Verify that the table blocks are stored in the buffer cache.

Answer: C

Explanation:

The problem statement clearly identifies at least one possible problem; that is the missing index on the columns in the join condition. If all queries were experiencing performance problems, then we might first verify that there are recent statistics on tables if the CBO is used. If objects in a specific tablespace were seen as performance constrained, we would check the file I/O statistics for the files in the tablespace.

Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 257:

Which dynamic performance view is the primary place to look for system performance and resource usage information?

- A. V\$SYSTEM_EVENT
- B. V\$PERFORMANCE
- C. V\$RESOURCE
- D. V\$SYSSTAT
- E. V\$SYSTEM_STATS

Answer: D

System performance and resource usage stats are kept in the V\$SYSSTAT view. Hit ratios, cumulative logons, total CPU usage, parse counts, total physical reads and writes, and in-memory and to-disk sorts are just a few of vital statistics that are available in this view. Wait events are recorded in the V\$system_event view; V\$PERFORMANCE, V\$RESOURCE, and V\$SYSTEM_STATS are fictions.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 258:

Which of these are true statements about tuning the dictionary cache? (Choose two.)

- A. The algorithm that manages data in the shared pool prefers to keep library cache data over dictionary cache data.
- B. The algorithm that manages data in the shared pool prefers to keep dictionary cache data over library cache data.
- C. When using sequences, the CACHE option improves dictionary cache performance.
- D. The DBA should perform DDL during peak activity hours.

Answer: B, C

It is important to remember that dictionary cache misses are more expensive than library cache misses, so the algorithm that manages the shared pool prefers to keep objects in the dictionary cache. If the library cache statistics look good, that means the library cache is sized appropriately and infers that the dictionary cache is sized appropriately. If you do not use the CACHE option on a sequence, each call to NEXTVAL creates a get in the dc_sequences dictionary cache item. Also, the DBA should perform DDL during off-peak hours.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 259:

Which of these accurately describes how the UGA is configured?

1Z0-033

- A. In both shared server and dedicated server environments, the session stack space is kept in the SGA.
- B. The sort areas and private SQL areas are included in the SGA in the dedicated server environment.
- C. In both shared server and dedicated server environments, the session stack space is kept in the PGA.
- D. The sort areas and private SQL areas are included in the PGA, regardless of the server configuration.

Answer: C

By definition, the session stack space is kept in the PGA, whether the process memory is assigned to a shared server or a dedicated server. The sort areas and private SQL areas are stored in the PGA for dedicated servers, and in the SGA for shared servers.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 260:

The default buffer cache can be expanded dynamically if needed, but only in granule-size increments of? (Choose two.)

- A. 2Mb if the SGA is less than 128Mb
- B. 4Mb if the SGA is less than 128Mb
- C. 8Mb if the SGA is less than 128Mb
- D. 8Mb if the SGA is larger than 128Mb
- E. 16Mb if the SGA is larger than 128Mb

Answer: B, E

The granule is defined as a unit of allocation in the dynamic SG

A. When an SGA component (DB buffer cache, shared pool) is expanded or shrunk, it must be in granule increments. The minimum number of granules allocated is 1 for each of the SGA structures (buffer cache, shared pool, and fixed SGA). The maximum number of granules is determined by MAX_SGA_SIZE, which is measured in bytes.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 261:

You want to allocate the memory structures in the shared pool so that you can begin collecting cache advice information at a later time. What is the correct setting for the DB_CACHE_ADVICE parameter?

- A. FALSE
- B. OFF
- C. READY

- D. ON
- E. ALLOCATE

Answer: C

Setting DB_CACHE_ADVICE to READY allocates the memory structures in the shared pool to collect the statistics needed to populate the V\$DB_CACHE_ADVICE view. The memory structures, and the view, are not populated with data until the DB_CACHE_ADVICE is set to ON, which signals the collection of buffer cache advice statistics. Setting the parameter to OFF stops collecting data and deallocates the shared pool memory structures.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 262:

Which of these tuning efforts should occur first?

- A. Tune the OS.
- B. Tune database I/O.
- C. Tune memory
- D. Tune the network.
- E. Tune the application.
- F. Tune the design.

Answer: F

Tuning the design is the first step in the application development phase of the Performance Tuning Model. Tuning the application should follow this step during the coding phase. The remaining tuning opportunities are usually the DBA's responsibility and occur later in the design and implementation process.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 263:

It appears that several user processes are at a standstill, that is, no throughput. However, other user processes seem to be working fine. Which dynamic performance view could you look at to ascertain the cause?

- A. V\$LIBRARYCACHE
- B. V\$LOCK
- C. V\$PARAMETER
- D. V\$INSTANCE
- E. V\$DB_OBJECT_CACHE

Answer: B

Explanation:

Of the choices, V\$LOCK should be the first place to look if some processes appear to be frozen while other processes are working without problems. If you see locked objects and sessions waiting on the locked objects, take action. Answers A and E are incorrect because they specifically show the state of the library cache and the objects cached there. Although these are good performance views to look at, they don't apply in this case. Answers C and D are incorrect; these views are informational only and won't help you resolve the locking conflict described in the question.
Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 264:

You have decided that a table is a good candidate for the RECYCLE pool. Which of these factors convinced you of that? (Choose two.)

- A. The table is small and used often.
- B. The table is greater than twice the size of the DEFAULT pool.
- C. The table is large and seldom used.
- D. The table is used often, but accessed solely with full table scans.
- E. The table is already CACHED.
- F. The table is less than 10 percent the size of the DEFAULT pool.

Answer: B, C

Recycle pool candidates have almost the exact opposite access pattern as KEEP pool candidates. A good candidate will be large and infrequently used. If the table is small and used often or the table is already CACHED it's more like a KEEP object. If the table is used often, but accessed solely with full table scans or the table is less than 10 percent the size of the DEFAULT pool it is good candidate for cached objects.
Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 265:

In the shared server environment, cursor state information resides in which pool? (Choose two.)

- A. The reserved shared pool
- B. The large pool if it is configured
- C. The PGA
- D. The shared pool if the large pool is not configured
- E. The large pool if the shared pool is not configured

Answer: B, D

Explanation:

In the shared server environment, UGA cursor state information and user session data are stored in the large pool if it is configured; otherwise, it is stored in the shared pool. This configuration uses less overall memory than the dedicated server environment.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 266:

Which of the following will not benefit from the NOLOGGING attribute? (Choice two.)

- A. CREATE INDEX
- B. Conventional path data loads
- C. Direct path data loads when the database is in noarchivelog mode
- D. CREATE TABLE ... AS SELECT
- E. Direct path data loads when the database is in archivelog mode

Answer: B, C

Explanation:

Conventional path data loads generate redo logging just like other Data Manipulation Language (DML) statement, so the NOLOGGING attribute has no effect. When the database is in noarchivelog mode, direct path data load do not generate redo; therefore the NOLOGGING attribute is irrelevant. To create index, create table as select from other table or to use direct path load with the database in archivelog mode you can benefit from the NOLOGGING attribute. CREATE INDEX and CREATE TABLE ... AS SELECT are two SQL statements that utilize NOLOGGING to improve performance.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 267:

Which of these is not an advantage of Oracle partitioning over manual striping? (Choose two.)

- A. Each partition can be assigned to its own tablespace and datafiles.
- B. Extent allocation is manual with the Oracle Partition option.
- C. The cost-based optimizer can choose better-performing execution plans by pruning partitions that aren't needed by the query.
- D. RAID-5 isn't compatible with manual striping, but it is compatible with partitioning.
- E. The administration of manual striping is labor-intensive.

Answer: B, D

Extent allocation and RAID-5 are not advantages of using the Oracle Partitioning option. Extent allocation is automatic with the Oracle Partitioning option, and the extents will end up in the partition's tablespace. You do not need to manually allocate an extent. RAID-5 is compatible with both manual striping and the Oracle Partitioning option.

Other choices are advantages of partitioning over manual striping.
Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 268:

Which of these SQL statements cause a sort operation? (Choose two.)

- A. CREATE TABLE
- B. GROUP BY
- C. UNION
- D. UNION ALL
- E. SORT BY

Answer: B, C

The GROUP BY clause explicitly directs Oracle to aggregate the result set, which requires the atomic data to be sorted. The UNION operator requires a sort because duplicate values are eliminated. CREATE TABLE is incorrect choice - the exception would be the CREATE TABLE ... AS SELECT when a primary key is specified. UNION ALL does not eliminate duplicates. SORT BY is not a valid SQL statement clause.
Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 269:

You discover contention for the cache buffers LRU chain latch. Which of the following describe a possible solution? (Choose two.)

- A. Increase the number of LRU latches on the buffer cache by increasing the LRU_CHAIN_LATCHES parameter value.
- B. Decrease the number of LRU latches on the buffer cache.
- C. Tune the application SQL to reduce large index range scans and full table scans.
- D. Increase the number of DBWn processes.
- E. Increase the number of shared server processes.
- F. Increase the number of ARCn processes.

Answer: C, D

Contention on the cache buffers LRU chain latch indicates that the buffer cache is very busy and has too much throughput. Investigate user SQL for full-table scans or inefficient index range scans. Also, consider tuning the DBWn processes and buffer cache settings. LRU_CHAIN_LATCHES is not a valid initialization parameter. Also we cannot directly control the number of cache buffers LRU chain latches. Increasing the number of shared server or ARCs processes are not relevant answers.
Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 270:

You're operating an OLTP database in manual undo management mode and notice that the query `SELECT COUNT(*) FROM V$TRANSACTION` is always less than or equal to 160 during peak transaction hours. How many rollback segments should you configure?

- A. 640
- B. 40
- C. 160
- D. 180
- E. 50
- F. 4
- G. 20

Answer: B

The general guideline is to configure the maximum number of concurrent transactions divided by 4. For this case, it's 40, so we should configure 40 rollback segments. Based on the guideline and the query, each of the other answers is an incorrect value.

Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 271:

You determine that a long-running transaction has used an unnecessarily high-level user-defined lock and is causing contention. Which of these courses of action will eliminate the contention problem without disrupting the transaction?

- A. Execute `DBMS_LOCK.allocate_unique()`
- B. Execute `DBMS_LOCK.release()`
- C. Execute `DBMS_LOCK.convert()`
- D. Execute `DBMS_LOCK.create()`

Answer: C

The `DBMS_LOCK.convert()` function converts a lock from one mode to another. The `kill()` function is nonexistent. `Allocate_unique()`, `create()` and `release()` are valid functions of the `DBMS_LOCK` package, but would not be used for this purpose. You may consider `release()` function, but that would completely release the user-defined lock.

Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 272:

You noticed that the high watermark for shared server process is continuously at the `MAX_SHARED_SERVER` limit. Choose a solution to optimize shared server performance:

1Z0-033

- A. Double the value for SHARED_SERVERS and MAX_SHARED_SERVERS.
- B. Increase MAX_SHARED_SERVERS by a few, and check instance performance.
- C. Increase SHARED_SERVERS by a few, and check instance performance.

Answer: B

Explanation:

One of the reasons we move to shared servers is to better utilize system resources. If you increase MAX_SHARED_SERVERS by too much, you may cause swapping or CPU starvation. If you increase SHARED_SERVERS too much, you may cause the same problems; remember, Oracle does not drop shared server processes below the number of SHARED_SERVERS. Increasing SHARED_SERVERS alone will not help if the shared server high watermark is the problem.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 273:

In a DSS application, the sales summary report queries data from seven different tables and groups the data by region and period. Which one of the following storage structures is a good design choice to improve query performance?

- A. IOT
- B. Cluster
- C. Materialized view

Answer: C

Explanation:

Materialized views are designed for the type of situation described in the question. A materialized view could be put in place to speed the queries that would have accessed multiple tables to produce summary information. A cluster would probably not be good structure because the described data is probably volatile and is probably not always queried together. An IOT is not a good way to store data from several tables; however, the materialized view could be stored as an IOT instead of a heap table.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 274:

You are building a large fact table for a data warehouse where there will be no updates or deletes. Which of the following settings for PCTFREE and PCTUSED makes the best performance combination?

- A. 0, 20
- B. 90, 10

- C. 10, 95
- D. 0, 40
- E. 20, 60
- F. 100, 0

Answer: A

A large fact table in a data warehouse is usually a table that is not updated. It is usually loaded, and then queried with full-table scans. Setting PCTFREE as slow as possible makes best sense, because there will be no updates into the table, and the table will be read with full-table scans. Combinations PCTFREE=10 and PCTUSED=95 or PCTFREE=100 and PCTUSED=0 are not possible. By setting PCTFREE to 90 and PCTUSED to 10 you will have result opposite of what we would expect, and dangerous because we would keep moving the block on and off the freelist with each inserted or deleted row. PCTFREE=0 and PCTUSED=40 is a acceptable answer, but there's really no need to set PCTUSED to anything since we won't drop below that threshold. PCTFREE=20 and PCTUSED=60 is a compromise - but not optimal.

Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 275:

Which of the following reasons to gather statistics with the DBMS_STATS package instead of the ANALYZE command? (Choose two.)

- A. The ANALYZE command doesn't support HISTOGRAMS.
- B. DBMS_STATS enables gathering statistics in parallel.
- C. The ANALYZE command does not compute statistics on indexes.
- D. DBMS_STATS enables block-level sampling.
- E. ANALYZE does not enable row-level sampling.

Answer: B, D

Explanation:

There are two advantages that DBMS_STATS has over the ANALYZE command. All other answers are incorrect because each is not true about the ANALYZE command; the ANALYZE command does create histograms, it does compute statistics on indexes, when specified, and it does enable row-level sampling. Although each of these is true, they are also available with DBMS_STATS.

Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 276:

Before a resource plan can become the active top-level plan, which of the following rules must be validated? (Choose all that apply.)

1Z0-033

- A. All referenced subplans and consumer groups must exist.
- B. The sum of percentage for resources that use the emphasis method cannot exceed 100 at any given priority level.
- C. The plan must have a directive for OTHER_GROUPS.
- D. All plans must contain at least one subplan.
- E. A plan has a maximum of 32 consumer groups and 32 subplans.

Answer: A, B, C, E

Each of the rules specified must be met for a plan to be validated. All referenced subplans and consumer groups must exist. CPU percentage for resources in the emphasis method cannot exceed 100 percent at any level, and there must be a directive for OTHER_GROUPS in the plan. There is a maximum of 32 consumer groups allowed in a plan, and a maximum of 32 subplans in a plan. Also no plan is required to have a subplan.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 277:

The purpose of a checkpoint is to? (Choose two.)

- A. Write the current redo log to the archived redo log.
- B. Synchronize the modified data blocks in memory with the datafiles on disk.
- C. Ensure database consistency.
- D. Ensure that data blocks are read into memory as quickly as possible.
- E. Keep track of redo entries.

Answer: B, C

Checkpoints are used to make sure that we have a consistent point to which we can recover in case of a system crash. To get there, we need to write the dirty blocks from the DB buffer cache to disk, and make sure that redo is available in the redo logs.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 278:

To reduce unnecessary sorting, you can employ which of the following SQL techniques? (Choose two.)

- A. Use the UNION operator instead of UNION ALL.
- B. Use DISTINCT when selecting the primary key.
- C. Don't use DISTINCT when selecting the primary key.
- D. Use the UNION ALL operator instead of UNION.
- E. Use the NOSORT option when building indexes on unsorted table data.

Answer: C, D

Explanation:

There's no need to use the DISTINCT keyword if you know you're selecting on a primary Key; the DISTINCT causes an unnecessary sort. Also, the UNION operator sorts the return set and removes duplicates; if you know there are no duplicates between the queries, then use UNION ALL instead.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 279:

Oracle uses latches for what purpose?

- A. To queue access to database objects such as tables
- B. To protect access to memory structures
- C. To keep memory doors closed
- D. To speed access to database tables
- E. To allow two or more process to access and modify the same structures at the same time

Answer: B

Explanation:

Latches are one of the methods used to protect Oracle structures; locks are the other one. Latches have no queuing feature; they not speed performance, nor do they allow simultaneous or shared access. Latches are exclusive and held briefly.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 280:

Which of these views would be used to diagnose tablespace performance problems?

- A. V\$TABLESPACE
- B. V\$FILESTAT
- C. V\$FILE_IO
- D. V\$TABLESPACE_USAGE
- E. DBA_TABLESPACES

Answer: B

The V\$FILESTAT view contains accumulated data about the number of physical reads and writes, blocks read and written, and the total read and write time. These statistics can be used to determine hot spots or disks that are much more active than others, and to evaluate the average read and write time for the files in a tablespace. The information can be used to determine if files need to be moved to different disk drives.

QUESTION 281:

Which of the following will reduce the number of disk sorts?

- A. Increasing SORT_AREA_SIZE
- B. Decreasing SORT_AREA_RETAINED_SIZE
- C. Increasing the SORT_MEMORY_ALLOCATED parameter
- D. Altering SORT_AREA_SIZE dynamically, based on available memory within SGA_MAX_SIZE

Answer: A

Explanation:

The value of SORT_AREA_SIZE determines the amount of memory that can be used for sorting operations; a lower value increase the likelihood that sorting will overflow the area and require disk extents. Increasing the size of the sort area will reduce free memory available to the remaining processes in the database server, but will also reduce the likelihood that sort operations will require disk space allocated from the temporary tablespace. SORT_AREA_SIZE, if too large, can contribute to paging and swapping.

Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 282:

The DBA can modify the default locking mechanism by setting which init.ora parameter?

- A. ROW_LOCKING=ROW
- B. ROW_LOCKING=TABLE
- C. ROW_LOCKING=DEFAULT
- D. ROW_LOCKING=INTENT

Answer: D

Explanation:

The DBA can modify the default locking mechanism by setting the ROW_LOCKING initialization parameter the default value is ALWAYS; this default locking at the row level during DML statements. This is the least restrictive. You can also specify DEFAULT, which is synonymous with ALWAYS. Setting the value to INTENT causes default locking at the table level, except for SELECT ... FOR UPDATE statements, which cause row-level locking.

Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 283:

What is PGA_AGGREGATE_TARGET?

- A. Initialization parameter for sort-based operators (for example, ORDER BY, GROUP BY, ROLLUP, window functions)
- B. Initialization parameter for hash-join operations
- C. Initialization parameter for bitmap merge and bitmap create operations
- D. Initialization parameter to enable automatic PGA memory management
- E. Initialization parameter to enable automatic PGA memory management if WORKAREA_SIZE_POLICY is set to MANUAL

Answer: D

Prior to Oracle9i, the maximum size of these working areas was controlled using the SORT_AREA_SIZE, HASH_AREA_SIZE, BITMAP_MERGE_AREA_SIZE and CREATE_BITMAP_AREA_SIZE parameters. Setting these parameters is difficult, because the maximum work area size is ideally selected based on the data input size and the total number of work areas active in the system. These two factors vary a lot from one work area to another and from one point in time to another. Thus, the various *_AREA_SIZE parameters are hard to tune under the best of circumstances. With Oracle9i, you can simplify and improve the way PGA memory is allocated, by enabling automatic PGA memory management. In this mode, Oracle dynamically adjusts the size of the portion of the PGA memory dedicated to work areas, based on an overall PGA memory target explicitly set by the DB

- A. To enable automatic PGA memory management, you have to set the initialization parameter PGA_AGGREGATE_TARGET, as described in the following section. Oracle 9i Performance Tuning Guide and Reference Oracle 9i Database Tuning with the Tuning Pack

QUESTION 284:

What should you do to import data and to avoid migration? (Choose two)

- A. Use the CASCADE option during import
- B. Use the COMPRESS option during export
- C. Use the RECORDLENGTH option during export
- D. Coalesce free space
- E. Add a new datafile

Answer: D, E

To avoid data migration during import you need coalesce free space or add a new datafile to tablespace. There is no CASCADE option for export/import utilities. Specifies the length, in bytes, of the file record. The RECORDLENGTH parameter is necessary when you must transfer the export file to another operating system that uses a different default value. The COMPRESS option specifies how Export and Import manage the initial extent

for table data.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 285:

After new version of OLTP applications implementation for testing several QA users have called to state that response time for applications is now very low.

You find the following:

1. The hit ratio in the buffer cache has fallen from its normal level of 88 to 90 percent and is only averaging 75 percent when you look at several different the reports from STATSPACK.
2. The GETHITRATIO for the library cache is 98%.
3. The disk sorts are rare and there is nothing suspicious in the datafiles read/writes section from STATSPACK.
4. Only one SQL query, specific for only one application, shows a huge amount of physical reads.

What should you do?

- A. Increase SORT_AREA_SIZE initialization parameter.
- B. Set WORKAREA_SIZE_POLICY parameter to AUTO and increase PGA_AGGREGATE_TARGET initialization parameter.
- C. Investigate, why there are so many physical reads for this application.
- D. Run DBMS_STATS package to compute tables/indexes statistics.

Answer: C

Since only one specific query shows extensive physical reads and all other wait events are normal you need first to tune this query. After that different SQL statement will be on the top of all time-consuming statements, which will require tuning also. Memory, disks I/O tuning should be considered only after SQL statements response time improvement. DBMS_STATS package can improve application response time, but it does not fix the core of the problem.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 286:

STATSPACK report shows a buffer busy waits. Detailed analysis shows that it's related with freelists contention. You need to fix this issue and add freelists.

- A. Search for a data block, causing contention and issue the command ALTER TABLE ... FREELIST.
- B. Search for a data block, causing contention and issue the command ALTER TABLESPACE ... FREELIST.
- C. Search for a segment header, causing contention and issue the command ALTER

TABLE ... FREELIST.

D. Search for a segment header, causing contention and issue the command ALTER TABLESPACE ... FREELIST.

Answer: C

Buffer busy wait event indicates that there are some buffers in the buffer cache that multiple processes are attempting to access concurrently. Query V\$WAITSTAT for the wait statistics for each class of buffer. Common buffer classes that have buffer busy waits include data block, segment header, undo header, and undo block. If the contention is on the segment header, then this is most likely freelist contention. A freelist is a list of free data blocks that usually includes blocks existing in a number of different extents within the segment. Blocks in freelists contain free space greater than PCTFREE. This is the percentage of a block to be reserved for updates to existing rows. In general, blocks included in process freelists for a database object must satisfy the PCTFREE and PCTUSED constraints. Specify the number of process freelists with the FREELISTS parameter. The default value of FREELISTS is one. The maximum value depends on the data block size. FREELISTS specifies the number of free lists for each of the free list groups for the schema object. Not valid for tablespaces.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 287:

Which description below describe OTHER_GROUPS consumer group?

- A. This consumer group cannot be explicitly assigned to a user. It must have a resource directive specified in the schema of any active plan. This group applies collectively to all sessions that belong to a consumer group that is not part of the currently active plan schema, including DEFAULT_CONSUMER_GROUP.
- B. This consumer group is the initial consumer group for all users/sessions that have not been explicitly assigned an initial consumer group. It has switch privileges granted to PUBLIC; therefore, all users are automatically granted switch privilege for this consumer group.
- C. It is the initial consumer group for the users SYS and SYSTEM
- D. It provides a group having lower priority than SYS_GROUP and OTHER_GROUPS in this plan.

Answer: A

Explanation:

This description fits the OTHER_GROUPS consumer group.

DEFAULT_CONSUMER_GROUP is the initial consumer group for all users/sessions that have not been explicitly assigned an initial consumer group.

DEFAULT_CONSUMER_GROUP has switch privileges granted to PUBLIC; therefore, all users are automatically granted switch privilege for this consumer group.

SYS_GROUP is the initial consumer group for the users SYS and SYSTEM.

LOW_GROUP provides a group having lower priority than SYS_GROUP and OTHER_GROUPS in this plan.

Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 288:

What can reduce performance with respect to redo logs? (Choose four)

- A. Place members of the same group on the same physical disk.
- B. RMAN slaves are too slow.
- C. Use RAID 5 for redo logs.
- D. Place redo logs on disks with temporary files.
- E. Impose table-level locks.

Answer: A, C, D, E

These actions can reduce redo logs performance and cause lock contention. Only slow RMAN slaves has nothing to do with redo logs performance reduction.

Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 289:

You have a parent table (DEPARTMENTS) and a child table (EMPLOYEES). Most of your queries use joins of these tables. You want to optimize access by foreign key between parent and child table. What should you do?

- A. Minimize space and time overhead for indexes associated with constraint.
- B. Set constraint deferred.
- C. Build NOT NULL constraint on foreign key.
- D. Build UNIQUE constraint on foreign key.
- E. Nothing, because the matching unique or primary key is never updated or deleted.

Answer: E

You should almost always index foreign keys. The only exception is when the matching unique or primary key is never updated or deleted.

Oracle 9i Performance Tuning Guide and Reference
Oracle 9i Database Tuning with the Tuning Pack

QUESTION 290:

You configured your database to use the Oracle Shared Server. Statistics shows that process S002 has 5 requests, S004 has 3 requests and S001, S006 and S008 has only 1 request. You also see that the process S008 needs 99.9% resources. What should you do?

- A. Make S002 a dedicated server.

- B. Make S008 a dedicated server.
- C. Make S004 a dedicated server.
- D. Make S001 a dedicated server.
- E. Make S006 a dedicated server.

Answer: B

In this case S008 process should be modified into a dedicated server, because it slows down all other shared servers and degrades Multithreaded server overall performance.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 291:

What answer below is the best for the number of latches Oracle will need:

- A. $(PHYRDS * PHYBLKRD + PHYWRTS * PHYBLKWRT) / 2$
- B. $(IMMEDIATE_GETS - IMMEDIATE_MISSES) / 2$
- C. $(IMMEDIATE_GETS - IMMEDIATE_MISSES + GETS - MISSES) / 4$
- D. Oracle will examine the values and automatically decide the best match for number of latches.

Answer: D

Do not tune latches. If you see latch contention, then it is a symptom of a part of SGA experiencing abnormal resource usage. Latches control access with certain assumptions (for example, a cursor is parsed once and executed many times). To fix the problem, examine the resource usage for the parts of SGA experiencing contention. Merely looking at V\$LATCH does not address the problem.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 292:

A question about what to do first in the tuning process?

- A. Tune the OS.
- B. Tune database I/O.
- C. Tune memory
- D. Tune the network.
- E. Tune the application.
- F. First monitor performance
- G. Tune the design.

Answer: F

You need to monitor database performance for a while before you start tuning process.

Oracle 9i Performance Tuning Guide and Reference

Oracle 9i Database Tuning with the Tuning Pack

QUESTION 293:

Your current transaction, an UPDATE of the employee table, is updating 1000 rows. What kind of locking mechanism is used?

- A. There is a shared lock on the table and a shared row lock on each row updated.
- B. All rows being updated hold a row-exclusive lock but no table locking is necessary.
- C. All rows being updated have a row-exclusive lock and an exclusive lock on the table.
- D. All rows updated have a row-exclusive lock on them and there is a shared lock on the table.

Answer: D

QUESTION 294:

After querying the V\$QUEUE table you note that there is an increasing number of items in the queue, and the total wait time is increasing. What should you do to improve performance of the database?

- A. Increase the size of the large pool by increasing the value of the initialization parameter LARGE_POOL.
- B. Increase the number of dispatchers available for use by the database, by increasing the value of the initialization parameter MAX_DISPATCHERS.
- C. Increase the number of shared servers available for use by the database by increasing the value of the initialization parameter MAX_SHARED_SERVERS.
- D. Increase the number of enqueue resources by increasing the value of the initialization parameter ENQUEUE_RESOURCES.
- E. Increase the number of job queue processes available to the database by increasing the value of the initialization parameter JOB_QUEUE_PROCESSES.

Answer: C

QUESTION 295:

Assuming you are given enough time and storage to reorganize poorly performing B-Tree indexes, which statement correctly describes the action you should take and the resulting benefit?

- A. Drop the indexes in question and create new indexes as bitmap indexes because bitmap indexes are compact.
- B. Drop the indexes in question and recreate them because it is always faster to recreate than to use any other method.
- C. Coalesce the indexes because it will move them to another tablespace and quickly free up index leaf blocks for use.
- D. Rebuild the indexes in question because that will shrink the heights of the trees wherever applicable.

Answer: D

QUESTION 296:

When tuning the dictionary cache, which init.ora parameter should you increase if there are too many cache misses due to the dictionary cache being incorrectly sized?

- A. DB_BLOCK_SIZE
- B. SORT_AREA_SIZE
- C. LARGE_POOL_SIZE
- D. SHARED_POOL_SIZE
- E. SHARED_POOL_RESERVE_SIZE

Answer: D

QUESTION 297:

Why are deadlocks uncommon in Oracle?

- A. Deadlocks cannot happen on a UNIX system.
- B. Oracle is very effective at escalating locks.
- C. Transactions explicitly override default Oracle locking.
- D. Oracle uses page-level locking instead of row-level locking.
- E. Oracle uses row-level locking instead of page-level locking.

Answer: E

QUESTION 298:

Application development and database configuration are closely related when DBAs need to tune database performance. Which statement correctly reflects a dependency that exists between the two areas?

- A. Highly normalized tables will provide optimal query performance.
- B. Uncommitted SQL statements are unlikely to affect database performance.
- C. Every type of SQL statement has the same impact on database performance.
- D. Tuning indexes may help avoid an excessive number of full table scans.
- E. The impact of adding a new application to the system is directly measurable by doubling the existing workload.

Answer: D

QUESTION 299:

You decide to improve the performance of the database buffer cache by creating

[1Z0-033](#)

multiple buffer pools. How do you create and define the buffer pools?

- A. Execute the buffer pool package.
- B. Set individual init.ora parameters for each buffer pool.
- C. Issue the ALTER SYSTEM CREATE butter pool <pool name> <size>; command.
- D. Issue the ALTER SYSTEM CREATE butter pool <pool name> <size>; command

Answer: B

QUESTION 300:

Which two statements about plan stability and stored outlines are true? (Choose two.)

- A. Stored outlines are saved in the OUTLN schema.
- B. You can have several stored outlines for a single SQL statement.
- C. You have to specify a category name when creating stored outlines.
- D. You can enable several stored outline categories at the same time.
- E. Stored outlines are recompiled when you analyze the associated objects.

Answer: A, B

QUESTION 301:

You determined that the value for REQUEST_FAILURES as seen from V\$SHARED_POOL_RESERVED is more than zero and always increasing. Which two actions could be appropriate? (Choose two.)

- A. Decrease the value for the LARGE_POOL_SIZE parameter.
- B. Increase the value for the LARGE_POOL_SIZE parameter.
- C. Increase the value for the SHARED_POOL_SIZE parameter.
- D. Decrease the value for the SHARED_POOL_SIZE parameter
- E. Increase the value for the SHARED_POOL_RESERVED_SIZE parameter.
- F. Decrease the value for the SHARED_POOL_RESERVED_SIZE parameter.

Answer: C, E

QUESTION 302:

You are using Oracle Shared Server. If the large pool is not configured, which two are stored in the shared pool instead of in private user memory? (Choose two.)

- A. bitmaps
- B. user sort results

- C. user session information
- D. cursor state information
- E. BLOBS

Answer: C, D

QUESTION 303:

Which two pieces of information are in the alert log and can be effectively used for database performance tuning? (Choose two.)

- A. process ID of PMON
- B. System Change Number
- C. non-default initialization parameters
- D. operating system authentication prefix
- E. statements that contain such keywords as ALTER, STARTUP, SHUTDOWN, ARCHIVE LOG, and RECOVER

Answer: C, E

QUESTION 304:

You work for a manufacturing company that frequently runs its production lines 24 hours a day. The accounting department is working with the development staff on plans to create an application to track manufacturing costs at the shop floor level. As a senior DBA, you have been assigned the task of ensuring that the new application will run with sub-second response time as information is entered into the application via terminals installed in the manufacturing area. Daily, weekly and monthly reports will also be run overnight detailing the previous period's activity. The reports must run during a time of high user interaction with the database. Choose two tasks that you should perform to make sure that the new application meets its sub-second response time requirement. (Choose two.)

- A. Help design the screen layout used to input the data into the application.
- B. Help with the design of the tables used in the application to guarantee that every table is normalized to the third normal form for good relational design.
- C. Use performance-monitoring tools to gather information about the current workload on the system and again later while testing the application to identify potential bottlenecks within the database.
- D. Help with the design of the tables used in the application to guarantee that every table is normalized to the third normal form for good relational design, and then de-normalize the tables as necessary to achieve the best performance of the application.

Answer: C, D

QUESTION 305:

A breakable parse lock _____.

- A. can cause waits
- B. can cause contention
- C. is held only by statements in the Library Cache
- D. is held only by PL/SQL objects in the Library Cache
- E. is held by statements and PL/SQL objects in the Library Cache

Answer: E

QUESTION 306:

Which three statements correctly describe UTLBSTAT/UTLESTAT and STATSPACK, and the major differences between them? (Choose three.)

- A. STATSPACK collects more data, including high-resource SQL, than UTLBSTAT/UTLESTAT.
- B. UTLBSTAT/UTLESTAT produces a report after collecting statistics while STATSPACK allows the report to be generated when desired.
- C. STATSPACK and UTLBSTAT/UTLESTAT will produce the same per transaction ratios given identical collection criteria and data.
- D. STATSPACK precalculates many useful ratios, such as cache hit ratios and per transaction and per second statistics, while many of these ratios must be calculated manually when using UTLBSTAT/UTLESTAT.
- E. STATSPACK is typically used for generating bare minimum statistical reports of database activity between 2 points in time, while UTLBSTAT/UTLESTAT extends the data capture to store all statistics in a database repository, allowing for better baseline setting and offline analysis.

Answer: A, B, D

QUESTION 307:

You want to run the UTLLOCKT script to detect lock contention. Which script must be run prior to running UTLLOCKT?

- A. CATPROC.sql
- B. CATLOCK.sql
- C. CATPERF.sql
- D. CATBLOCK.sql
- E. CATBPROC.sql

Answer: D

QUESTION 308:

What are the steps the server process goes through when it needs to access a particular data block into the buffer cache?

- A. The server process must always read the current version of a data block from the data files into the buffer cache. Once it is there, the server process can access the data block.
- B. The server process always contacts the DBWR process to read the current version of a data block from the data files into the buffer cache. Once it is there, the server process can access the data block.
- C. The server process first checks to see if the block is already in the buffer cache using a hash function. If it is already in memory, the server process uses it. If the block is not in memory, the server process reads the data block into the buffer cache. Once it is there, the server process can access the data block.
- D. The server process first checks if the block is already in the buffer cache using a hash function. If it is already in memory, the server process uses it. If the block is not in memory, the server process contacts the DBWR process to read the data block into the buffer cache. Once it is there, the server process can access the data block.

Answer: C

QUESTION 309:

Database Resource Manager uses resource plans to determine resource limits for the set of users. Which statement is true?

- A. Resource plans are set using profiles.
- B. Only one resource plan can be stored in the database at one time.
- C. The database can have many resource plans, but only one can be active at any one time.
- D. The database can have many resource plans, and each user chooses which plan to belong to.

Answer: C

QUESTION 310:

You have free buffer waits occurring in your Oracle9 i database. You determine that DBWn performance is the problem. Which two situations could cause the problem? (Choose two.)

- A. The I/O system is too slow.
- B. The buffer cache is too large.
- C. DBWn is waiting for resources such as latches.
- D. The DBWn process is busy satisfying users' read requests.
- E. Buffers are being kept in the OS cache, instead of going straight to disk.

Answer: A, C

QUESTION 311:

The buffer cache advisory is used for which two purposes? (Choose two.)

- A. Provides a physical read factor for the keep buffer pool if it is in use in the SGA.
- B. Advises DBAs on the optimal size of the default buffer cache for the current workload.
- C. Advises DBAs on whether to set up the keep and recycle buffer pools if neither is defined in the SGA.
- D. Advises DBAs on whether to set up buffer caches with different block sizes when only the default buffer cache is being used in the SGA.

Answer: A, B

QUESTION 312:

You want to gather statistics for all objects in a schema, but you don't want to discard the current statistics. What is the most efficient method of completing this task with minimal impact on database downtime and performance?

- A. Create a clone database and gather statistics from it.
- B. Take DB backup, gather statistics, then restore the database.
- C. Export tables in the schema, gather statistics, then import the tables.
- D. Use the DBMS_STATS package to gather statistics while saving the current statistics, then restore the original statistics.

Answer: D

QUESTION 313:

Which dynamic performance view should you query to find out how many sorts went to disk and how many sorts went to memory?

- A. V\$LOG
- B. V\$CACHE
- C. V\$SYSSTAT
- D. V\$SESSION
- E. V\$SQLAREA
- F. V\$ROWCACHE

Answer: C

QUESTION 314:

You and several other DBAs are discussing the different areas in the SGA and debating the importance of tuning each area. The four statements below are made during your conversation. Which two are true? (Choose two.)

- A. When you tune the shared pool, you should concentrate on tuning the data dictionary cache. If the hit ratio for the data dictionary cache is acceptable, the hit ratio for the library cache should also be acceptable.
- B. A cache miss in the database buffer cache is more expensive than a miss in the data dictionary cache or the library cache, so tuning the database buffer cache is a priority.
- C. A cache miss in the data dictionary cache or the library cache is more expensive than a miss on the database buffer cache, so tuning the shared pool is a priority.
- D. When using shared servers, you should either allocate additional space to the shared pool or set up the large pool to accommodate the User Global Area for the shared server users.

Answer: C, D

QUESTION 315:

While tuning memory for an Oracle9i database, you executed this query:

Exhibit:

```
SELECT sid, event, seconds_in_wait, state
FROM v$session_wait
WHERE name LIKE 'log buffer%';
```

SID	EVENT	SECONDS_IN_WAIT	STATE
5	log buffer space	110	WAITING

Which two statements are true? (Choose two.)

- A. You should move the log files to faster disks.
- B. You should make the log buffer bigger.
- C. You should perform inserts and updates with the NOLOGGING attribute to reduce redo.
- D. You should use nologging when doing conventional direct load into a database in noarchivelog mode.

Answer: A, B

QUESTION 316:

You want to know the number of deleted values for one of your indexes, so you analyzed the index with the VALIDATE STRUCTURE option. What should you do now?

- A. Query the INDEX_STATS view.
- B. Analyze the index with the COMPUTE or ESTIMATE STATISTICS option.
- C. Query DBA_HISTOGRAMS to investigate the spread of rows across index key values.
- D. Query the DBA_INDEXES view, and specify a WHERE clause to select the right index.

Answer: A

QUESTION 317:

When is the best time to PIN large or frequently-required objects in the SHARED POOL?

- A. when memory becomes available
- B. at the start of each user session
- C. the first time the object is called
- D. at database startup before loading any objects into memory
- E. at instance startup; load the object into memory and pin it

Answer: E

QUESTION 318:

Which statement is true when connections to the Oracle Instance use the shared server configuration?

- A. The User Global Area (UGA) resides in the private user memory.
- B. The User Global Area (UGA) contains only the user session data.
- C. The User Global Area (UGA) contains only the cursor state information.
- D. The User Global Area (UGA) resides in the shared pool if the large pool is not defined.

Answer: D

QUESTION 319:

Repeatedly deleting large number of rows from tables over a period of time will result in massively wasted extents because unused space below each High-Water Mark is never released. You are unable to issue TRUNCATE <tablename>; DROP STORAGE to reset the High-Water Mark because you need to sustain some rows in the tables. What should you do to release the unused space?

- A. Execute DBMS_SPACE.UNUSED_SPACE procedure and delete the file identified by the LAST_USED_EXTENT_FILE_ID.

[1Z0-033](#)

- B. Issue ALTER TABLE <tablename> MOVE to move tables in question to healthy tablespace(s).
- C. Issue DROP TABLE <tablename> to remove storage characteristics stored in the data dictionary so that old High-Water Marks recorded in the segment headers will be eliminated.
- D. Issue ANALYZE TABLE <tablename> COMPUTE STATISTICS and delete the rows from the table if the value returned in the CHAIN_CNT column is more than 0.
- E. Issue ALTER TABLE <tablename> DEALLOCATE UNUSED to release unused space below High-Water Mark and repeat it until all of the tables in question are taken care of.

Answer: B

QUESTION 320:

New Resource Manager consumer groups, plans, or directives are stored in the pending area until they are _____.

- A. validated
- B. dropped from the database
- C. assigned to a database role
- D. assigned to a database user
- E. recorded in the archive logs

Answer: A

QUESTION 321:

Which of these is NOT in the shared pool?

- A. Object Cache
- B. Library Cache
- C. Shared SQL/PLSQL
- D. Data Dictionary Cache

Answer: A

QUESTION 322:

You created the COMPANY table with the NOLOGGING attribute:
Exhibit:

[1Z0-033](#)

```
CREATE TABLE company (  
  company_id          NUMBER(5) not null,  
  company_name        VARCHAR2(4) NOT NULL,  
  address              VARCHAR2(200)  
)  
NOLOGGING  
storage (initial 32K next 32K pctincrease 0)  
;
```

Which two operations on the COMPANY table will suppress redo entry generation?
(Choose two.)

- A. DELETE FROM COMPANY;
- B. Direct Path loads using SQL*Loader
- C. DELETE FROM COMPANY WHERE company_id = 100;
- D. Direct-load INSERT operation on the COMPANY table
- E. UPDATE COMPANY SET COMPANY_ADDRESS = 'One Oracle Way' WHERE company_id = 100

Answer: B, D

QUESTION 323:

Which two statements about row migration are true? (Choose two.)

- A. Row migration is caused by a PCTFREE value set too low.
- B. Row migration can be resolved using the ANALYZE command.
- C. Row migration can be reduced by choosing a larger block size.
- D. Queries that use an index to select migrated rows perform additional I/O.

Answer: A, D

QUESTION 324:

If segments have too few free lists, your database performance may suffer because of contention for _____.

- A. free list latches
- B. segment header blocks
- C. free list entries in data blocks
- D. transaction slots in rollback segments

Answer: B

QUESTION 325:

As a conscientious DBA you are aware that tuning should begin at the initial stage of the application development phase. Assuming you have optimal data model and

1Z0-033

data file layout to eliminate hot spots, what is the next step you should take to isolate the problems that may cause poor database performance?

- A. tune indexes
- B. increase SGA
- C. add more CPUs
- D. add more LRU latches
- E. monitor performance

Answer: E

QUESTION 326:

What does this statement accomplish? SQL> ANALYZE INDEX index_name VALIDATE STRUCTURE;

- A. It provides information in the INDEX_HISTOGRAM view to indicate whether an index is invalid or valid.
- B. It provides information in the DBA_INDEXES view for the COST BASED Optimizer when choosing an execution plan.
- C. It places information into the INDEX_STATS view that provides statistics on the space used by an index.
- D. It checks to see if each index entry points to a valid row in the underlying table and saves the results in the DBA_INDEXES view.
- E. It checks to see if each row in the underlying table has an index entry and saves the results into the DBA_INDEXES view.

Answer: C