

1z0-033

TEST KING



LEADING THE WAY IN IT
TESTING AND CERTIFICATION TOOLS!

Oracle9i Performance Tuning Study Guide

Version 3.1

Important Note

Please Read Carefully

Study Tips

This product will provide you questions and answers along with detailed explanations carefully compiled and written by our experts. Try to understand the concepts behind the questions instead of cramming the questions. Go through the entire document at least twice so that you make sure that you are not missing anything.

Latest Version

We are constantly reviewing our products. New material is added and old material is revised. Free updates are available for 90 days after the purchase. You should check for an update 3-4 days before you have scheduled the exam.

Here is the procedure to get the latest version:

1. Go to www.testking.com
2. Click on **Login** (upper right corner)
3. Enter e-mail and password
4. The latest versions of all purchased products are downloadable from here. Just click the links.
Note: If you have network connectivity problems it could be better to right-click on the link and choose **Save target as**. You would then be able to watch the download progress.

For most updates it enough just to print the new questions at the end of the new version, not the whole document.

Feedback

Feedback on specific questions should be send to feedback@testking.com. You should state

1. Exam number and version.
2. Question number.
3. Order number and login ID.

We will answer your mail promptly.

Copyright

Each pdf file contains a unique serial number associated with your particular name and contact information for security purposes. So if you find out that particular pdf file being distributed by you. Testking will reserve the right to take legal action against you according to the International Copyright Law. So don't distribute this PDF file.

Note: There will be a major update of this product later in October. We recommend that you wait until this update before scheduling the exam.

QUESTION NO: 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 NO: 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 NO: 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 NO: 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 NO: 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:

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 NO: 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 NO: 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 NO: 8

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

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

Answer: A

Explanation:

GETS is number of times a Willing-to_Wait latch was acquired without waiting. So GETS statistic will be incremented.

Incorrect Answers

- B:** SLEEPS statistic provides the number of times a process had to wait before obtaining a Willing-toWait latch.
- C:** MISSES shows the number of times a Willing-to_Wait latch was not acquired and a wait resulted.
- D:** IMMEDIATE_GETS is statistic of numbers of times an Immediate latch was acquired without waiting.
- E:** IMMEDIATE_MISSES is the number of times an Immediate latch was not acquired and a retry resulted.

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

QUESTION NO: 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 NO: 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 NO: 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.

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.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 150
Chapter 3: SQL Application Tuning and Design

QUESTION NO: 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 NO: 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 NO: 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 NO: 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\$BYUFFER_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 NO: 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 NO: 17

In which two ways can you reduce the amount of sorting that is performed? (Choose two)

- A. By using UNION instead of UNION ALL.
- B. By using NOSORT when creating tables.
- C. By using NOSORT when creating indexes.
- D. By using COMPUTE instead of ESTIMATE when analyzing objects.
- E. By reducing the number of users that have the sort privilege.

F. By creating appropriate indexes on tables that are joined often.

Answer: B, F

Explanation:

You can avoid a sorting operation by using NOSORT clause when creating table or by creating appropriate indexes on tables that are joined often.

Incorrect Answers

A: The UNION type SQL statement will cause sort operation.

C: You cannot avoid sorting by using NOSORT when creating indexes.

D: By using COMPUTE instead of ESTIMATE when analyzing objects you will process ALL records in the table. It will cause sorting also.

E: There is no sort privilege in Oracle. All users can sort data in Oracle tables.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 411-412
Chapter 8: Tuning Disk I/O

QUESTION NO: 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.

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

Chapter 8: Tuning Disk I/O

QUESTION NO: 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 NO: 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.

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

QUESTION NO: 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 NO: 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 NO: 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 NO: 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 DBA. 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 NO: 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 NO: 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 NO: 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 NO: 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 Sever 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 NO: 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, Jaseph C. Johnson, p. 517-547

Chapter 10: Operating System Tuning

QUESTION NO: 30

Which three actions will cause queries to place a table's blocks at the most-recently-used end of the LRU list? (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, C, 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. Blocks will be placed at the most-recently-used end of the LRU list if the query performs a full table scan.

Incorrect Answers

- B:** Querying the table by using a CACHE hint will not cause queries to place a table's blocks at the most-recently-used end of the LRU list.
- 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 NO: 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 NO: 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 NO: 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 NO: 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 NO: 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 NO: 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 NO: 37**Which two statements about row migration are true? (Choose two)**

- 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 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. Setting PCTFREE more high can fix this issue.
- 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 NO: 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 NO: 39

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

- A. Memory
- B. Data dictionary.
- C. Temporary tablespace.
- D. Corresponding tablespace itself.

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 NO: 40

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

- A. In Oracle8i, 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 date-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 NO: 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 NO: 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 NO: 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 NO: 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 NO: 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 NO: 46

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

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:

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 NO: 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 NO: 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.
- 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 NO: 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 NO: 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 NO: 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 NO: 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 NO: 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.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 182-186
Chapter 3: SQL Application Tuning and Design

QUESTION NO: 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 populates 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 NO: 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 NO: 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 even when the materialized view's data has not been refreshed and is inconsistent with the underlying detail data in the base tables.

Answer: C

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 NO: 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.
- 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 NO: 58

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 locally managed tablespace if tables are stored in locally managed tablespaces.

Incorrect Answers

- A:** This information is stored in the corresponding locally managed tablespace 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 NO: 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 NO: 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 NO: 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 NO: 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 NO: 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 NO: 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, C, D

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

Incorrect Answers

B: It may not indicate disk I/O contention on the redo log files.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 329-330
Chapter 7: Tuning Redo Mechanisms

QUESTION NO: 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 NO: 66

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

- A. More concurrency.
- B. Faster full table scans.
- C. Fast primary key-bases access.
- D. Less contention on the segment header.
- E. No duplication of primary key values storage.

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 NO: 67

Which view shows the number of full table scan?

- 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 NO: 68

Which two views would you query to monitor sessions relates 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 NO: 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 NO: 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 NO: 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 NO: 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 NO: 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 NO: 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 NO: 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 NO: 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 NO: 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 NO: 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 NO: 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 NO: 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 NO: 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 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 NO: 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.
- 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 NO: 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 NO: 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 NO: 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, F

Explanation:

To avoid of increasing the value for REQUEST_FAILURES you should increase the value for SHARED_POOL_SIZE parameter and decrease 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.
- E: You should decrease, not increase, 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 NO: 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 NO: 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:

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 NO: 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 NO: 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 NO: 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 NO: 91

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

- 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 NO: 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. ALTER SYSTEM FLUSH SHARED_POOL;
- B. EXECUTE dbms_shared_pool.unpin;
- C. EXECUTE dbms_shared_pool.unkeep;
- D. EXECUTE dbms_library_cache.unpin;

Answer: C

Explanation:

Too 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 NO: 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 NO: 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 NO: 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 NO: 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 the 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 NO: 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 NO: 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 NO: 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.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 338-340
Chapter 7: Tuning Redo Mechanisms

QUESTION NO: 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 NO: 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 NO: 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 NO: 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 NO: 104

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

The database has 3 Redo Log Groups with 1 member each of size 2 MB.

The LOG_CHECKPOINT_INTERVAL parameter is set to 10000.

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 NO: 105

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 then 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 NO: 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 NO: 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 NO: 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 form 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 form 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 NO: 109

These are the current settings for your Oracle database:

The database has 3 Redo Log Groups with 1 member each of size 2 MB.

The LOG_CHECKPOINT_INTERVAL parameter is set to 10000.

The LOG_CHECKPOINT_TIMEOUT parameter is set to 1800.

**Which action(s) would increase the number of checkpoints in your database? (Choose all that apply.)
Assume you will restart the database afterwards.**

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

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 from 1800 to 2400 you will decrease the database checkpoints.
- B: Changing the LOG_CHECKPOINT_INTERVAL parameter from 10000 to 12000 you will decrease the database checkpoints.
- C: Increasing the size of all the Redo Log Group member files to 3 MB will decrease the database checkpoints.

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

QUESTION NO: 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 NO: 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 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 NO: 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 NO: 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 NO: 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 then 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 NO: 115

A latch is defined 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 defined 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 NO: 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 NO: 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 SGA. 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.

OCP: Oracle 9i Performance Tuning Study Guide, Joseph C. Johnson, p. 472-475
Chapter 9: Tuning Contention

QUESTION NO: 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.

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

QUESTION NO: 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 NO: 120

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

- A. Avoid row chaining and migration.
- B. Keep the b-tree structure densely clustered.
- C. Speed up full table scans and fast full index scans.
- D. Improve performance when the index-organized table is clustered.

Answer: B

Explanation:

You need to use a row overflow to keep the B-Tree index structure densely clustered.

Incorrect Answers

- A: This structure is not used to avoid row chaining and migration.
- C: IOTs are not good candidates to improve full table scans and fast full index scans. Row overflow is used for different reason.
- D: It does not improve performance when the index-organized table is clustered.

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

QUESTION NO: 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 NO: 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 NO: 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, DBW_n 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, DBW_n 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 NO: 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 NO: 125

Sometimes the LGWR process must wait because DBW has not compresses 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 to 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 NO: 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 NO: 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 NO: 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 NO: 129

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

- A. You did not refresh snapshots in time.

- B. An ITL entry in a data block has been reused.
- C. 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 NO: 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 NO: 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 NO: 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