

Oracle9i Database: Performance Tuning

1Z0-033

This study guide consists of 70 questions with answers and explanations.

Question 1.

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.

Option B:

Long running serializable transactions do not require large rollback segments.

Option C:

Reports do not modify any table records. So they do not need specific rollback segment.

Option D:

Discrete transactions that modify many rows in the same block will not cause rollback segment contention.

Question 2.

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.

Option A:

There are at least 50, not 5, blocks per LRU latch for each pool.

Option B:

There are at least 50, not 10, blocks per LRU latch for each pool.

Option C:

There are at least 50, not 30, blocks per LRU latch for each pool.

Option E:

There are at least 50, not 100, blocks per LRU latch for each pool..

Question 3.

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.

Option D:

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

Question 4.

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 non default parameters when the instance is started.

Option B:

Online redo log is not used to store this information.

Option C:

Archive redo log is just archived copy of online redo log. It does not have any information about non default parameters.

Option D:

SYSTEM's user's trace file is not used to provide this information.

Question 5.

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.

Option 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.

Option D:

Transactions should never wait for access to rollback segment.

Question 6.

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.

Option A:

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

Option B:

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

Option C:

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

Question 7.

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.

Option A:

You need to use as less memory as possible.

Option B:

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

Option C:

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

Question 8.

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.

Option A:

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

Option B:

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

Option 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.

Question 9.

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.

Option B:

Large amount of rollback segments decrease possibility of lock contention.

Option C:

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

Option D:

Lock contention is not related with shared pool size.

Question 10.

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.

Option 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.

Option C:

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

Option D:

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

Question 11.

Database Resource Manager uses resource plans to determine resource limits for the set of users.

Which statement is true in reference to resource plans?

- 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.

Option A:

Resource plans are not set using profiles.

Option B:

The database can have many resources plans.

Option D:

Each user cannot choose which plan to belong to because only one plan can be active at any one time.

Question 12.

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.

Option 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.

Option D:

Defining the table without the option for caching will not place blocks in the LRU list.

Option F:

This statement will not cause queries to place a table's blocks at the most-recently-used end of the LRU list.

Option G:

You don't need to create a separate database buffer cache to hold cached table.

Question 13.

Which four activities cause a sort? (Choose four)

- A. Creation of an index
- B. Execution of the ANALYZE command
- C. Use of the ORDER BY and GROUP BY clauses
- D. Use of insert operations involving index maintenance.
- A. Use of the UNION, INTERSECT AND MINUS operations.

Answer: A, C, D & E

Explanation:

All these activities can cause a sort operation. Only if you are creating an index with NOSORT option sorting will not have place.

Execution of the ANALYZE command does not require sorting.

Question 14.

The Accounting department has contacted you concerning their perceived poor performance of the new accounting package that was installed last week. Which three actions would help diagnose the problem? (Choose three)

- A. Query the DBA_ROLLBACK_SEGS table to see if there is contention for rollback segment headers while the accounting application is being run.
- B. Query DBA_DATA_FILES while the accounting application is being used to determine if any of the data files are being used excessively.
- C. Query the V\$SQL_PLAN table to determine which, if any, indexes are being used for the accounting tables.
- D. Query the V\$FILESTAT table while the accounting application is being used to determine if any of the data files are being used excessively.
- E. Query the V\$LOGFILE table to see if an excessive amount of redo is being generated while the accounting application is being run.
- F. Generate performance monitoring reports using UTLBSTAT and UTLESTAT or the STATSPACK while the account application is being used and compare them to your previously generated base line reports from the same tool.

Answer: C, D & F

Explanation:

The V\$SQL_PLAN view provides a way of examining the execution plan for cursors that were executed and are still cached. The V\$FILESTAT view contains detailed file I/O statistics for each file, including the number of I/Os for each file and the average read time. Generate performance monitoring reports using UTLBSTAT and UTLESTAT or the STATSPACK will also help to find the reason of poor performance.

Option A:

The DBA_ROLLBACK_SEGS data dictionary view stores information about the rollback segments of a database. But this information cannot be used to check contention for rollback segment headers.

Option B

The DBA_DATA_FILES data dictionary view provides descriptive information about each datafile. including the tablespace to which it belongs and the file id. The file id can be used to join with other views for detail information, but itself it will not give you information if any of the data files are being used excessively.

Option E:

This view just identifies redo log groups and members and member status.

Question 15.

Which three are used when administrating systems with the database Resource Manager? (Choose three)

- A. Histograms
- B. Explain Plan
- C. Resource Plan
- D. Resource User Group
- E. Resource Consumer Group
- F. Resource Plan Directives
- G. Resource Plan Assignments

Answer: C, E & F

Explanation:

Resource Plan, Resource Consumer Group and Resource Plan Directives are used when administering systems with the database Resource Manager.

Option A:

Histograms are not used with the database Resource Manager

Option B:

Explain Plan has no relation with the database Resource Manager

Option D:

There is Resource Consumer Group, not Resource User Group in Resource Manager

Option G:

There is not Resource Plan Assignments for Resource Manager

Question 16.

There are two users, John and Susan, who are updating the EMPLOYEE table in the following order:

1. First, John issues this SQL update:

```
UPDATE EMPLOYEE SET LAST_NAME = 'SMITH' where ID=200;
```

2. Next, Susan Issues this SQL update:

```
UPDATE EMPLOYEE SET SALARY = 5000 WHERE ID=250'
```

3. Next, John issues this SQL update:

```
UPDATE EMPLOYEE SET LAST_NAME 'BAKER' WHERE ID=250
```

4. FINALLY, Susan issues this SQL update:

```
UPDATE EMPLOYEE SET SALARY=60000 WHERE ID=200;
```

What will be the result?

- A. Oracle kills Susan's session to prevent a deadlock.
- B. Oracle will detect a deadlock and roll back Susan's session
- C. Oracle kills both John's and Susan's statements to prevent a deadlock
- D. Both John's and Susan's sessions would hang indefinitely because of a deadlock
- E. Oracle will detect a deadlock and roll back the statement causing the deadlock.

Answer: B

Explanation:

Oracle will detect a deadlock and roll back the Susan's session : session detected a deadlock.

Option A:

Oracle will detect a deadlock FIRST and roll back the statement detected a deadlock. It will be Susan's session.

Option C:

Oracle will not kill both statements, only one, detecting a deadlock.

Option D:

Both John's and Susan's sessions would not hang indefinitely: Oracle will handle a deadlock in most cases.

Option E:

Oracle will detect a deadlock and roll back the statement DETECTED the deadlock, not causing it.

Question 17.

Which two statements are valid settings for the initialization parameter 'USE_STORED_OUTLINES'? (Choose two)

- A. SQL> alter session set USE_STORED_OUTLINES = TRUE;
- B. SQL> alter session set USE_STORED_OUTLINES =PUBLIC;
- C. SQL> alter session set USE_STORED_OUTLINES = PRIVATE;
- D. SQL> alter session set USE_STORED_OUTLINES = DEFAULT;
- E. SQL> alter session set USE_STORED_OUTLINES = <outline name>;
- F. SQL> alter session set USE_STORED_OUTLINES = <category name>;

Answer: A & F

Explanation:

There are two valid settings for the initialization parameter USE_STORED_OUTLINES: TRUE or category name.

Option B:

This statement is incorrect

Option C

This statement is incorrect

Option D:

This statement is incorrect

Option E:

This initialization can use category name, not outline name.

Question 18.

Which two statements are true regarding the use of DB_CACHE_ADVICE init.ora parameter?(Choose two)

- A. Setting the parameter to READY reserves space in the buffer cache to store information about different buffer cache sizes, but no CPU overhead is incurred
- B. Setting the parameter to READY reserves space in the shared pool to store information about different buffer cache sizes, but no CPU overhead is incurred.
- C. Setting the parameter to ON reserves space in the buffer cache to store information about different buffer cache sizes, and CPU overhead is incurred as statistics are collected
- D. The V\$DB_CACHE_ADVICE view contains information that predicts the estimated number of physical reads for different cache sizes for each buffer cache setup in the SGA.

Answer: C & D

Explanation:

Each cache size simulated has its own row in this view, with the predicted physical I/O activity that would take place for that size. The **DB_CACHE_ADVICE** parameter is dynamic, so dynamic, so the advisory can be enabled and disabled dynamically to allow you to collect advisory data for a specific workload.

Two minor overheads are associated with this advisory:

- CPU: When the advisory is enabled, there is a small increase in CPU usage, because additional bookkeeping is required.
- Memory: The advisory requires memory to be allocated from the shared pool (about 100 bytes for each buffer)

The `V$DB_CACHE_ADVICE` view is populated when the **DB_CACHE_ADVICE** parameter is set to **ON**. This view shows the simulated miss rates for a range of potential buffer cache sizes.

Option A:

The `DB_CACHE_ADVICE` `init.ora` parameter should be set to `ON`, not `READY`.

Option B

The `DB_CACHE_ADVICE` `init.ora` parameter should be set `ON`, not `READY`. Also it reserves space in the buffer cache, not the shared pool.

Question 19.

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

- A. Full table scans are generally faster on cluster tables than on non-clustered tables.
- B. Hash clusters may be desirable when the number of key values is predictable and key values are evenly distributed.
- C. If tables in a hash cluster require more space than the initial allocation for the cluster, performance degradation can be substantial because overflow blocks are required.
- D. Storing a single table in a hash cluster can be useful regardless of whether the table is joined frequently with other tables or not, provided other criteria for choosing a hash cluster are met.

Answer: A, B & C

Explanation:

These statements are true with respect to clusters.

Option B:

Storing a single table in a hash cluster cannot be useful in some cases. It's better to use clusters if the table is joined frequently with other tables.

Question 20.

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

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

Answer: D

Explanation:

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

Option A:

`CATPROC.sql` script is required to load the initial PL/SQL environment.

Option B:

There is no CATLOCK.sql script in Oracle9i

Option C:

There is not CATPERF.sql script in Oracle9i

Option E:

There is no CATBPROC.sql script in Oracle9i

Question 21.

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

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

Answer: D

Explanation:

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

Option A:

PGA will not be used for this purpose

Option B:

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

Option C:

LARGE POOL is not configured as question states

Option E:

There is no SHARED SQL AREA in Oracle9i

Option F:

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

Question 22.

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..

Option B:

You can have as many stored outline per SQL statement as you need.

Option D:

Stored outlines are saved in the OUTLN schema.

Option E:

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

Question 23.

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.

Option B:

This command does not provide information in the INDEX_HISTOGRAM view.

Option C:

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

Question 24.

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.

Option A:

Locking statistics are not reported in REPORT.TXT file.

Option B:

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

Option C:

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

Question 25.

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: A, C

Explanation:

Shared code and bind variable should be used because quick parsing is required and due to high level of DML activity on the OLTP systems you need to use more small rollback segments rather than a few large rollback segment and dynamic space allocation is expensive so build the data object with appropriate size.

Question 26.

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.

Option A:

There is no TEMP tablespace in Oracle by default.

Option B:

There is no DATA tablespace in Oracle by default.

Option D:

ROLLBACK tablespace cannot be used as temporary tablespace.

Question 27.

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.

Option 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.

Option B:

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

Option 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.

Question 28.

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 if 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.

Option 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.

Question 29.

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.

Option A:

The UNION type SQL statement will cause sort operation.

Option C:

You cannot avoid sorting by using NOSORT when creating indexes.

Option D:

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

Option E:

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

Question 30.

What will this statement do?

```
CREATE TABLESPACE temp
```

```
DATAFILE 'C:\database\temp.dbf' SIZE 10m
```

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.

Option 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.

Option B:

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

Option D:

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

Question 31.

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

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

Answer: A

Explanation:

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

Option B:

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

Option 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.

Option D:

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

Question 32.

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

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

Answer: C

Explanation:

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

Option A:

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

Option B:

Large table cannot be easily 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.

Option D:

Large table is not good candidate to be cached.

Question 33.

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

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

Answer: D

Explanation:

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

Option A:

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

Option B:

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

Option C:

LOG_ARCHIVE_DEST shows the directory for the archived redo logs.

Question 34.

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

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

Answer: B, C & E

Explanation:

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

Option 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.

Option 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.

Question 35.

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

- A. DSS environment prefer a large block size..
- B. Small block sizes result in more block contention.
- C. Random access to large object favors 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.

Option B:

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

Option C:

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

Option 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.

Question 36.

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

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

Answer: D

Explanation:

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

Option A:

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

Option B:

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

Option C:

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

Question 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.

Option C:

Row migration cannot be reduced by choosing a larger block size. Setting PCTFREE more high can fix this issue.

Option 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.

Option E:

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

Question 38.

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

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

Answer: A, B & E

Explanation:

The ROLLSTAT view contains detailed information regarding the behavior of the rollback segments in the database. In particular, the columns USN, GETS and WAITS are particularly useful for measuring contention for the rollback segment's header. WAITS column shows number of times a user Server Process needed to access the rollback segment header and experienced a wait. The V\$WAITSTAT view contains information on block contention statistics. A nonzero value in the UNDO_HEADER column indicates a contention problem. The V\$SYSTEM_EVENT view tracks performance related information on rollback segments via the Undo Segment TX Slot statistic. Ideally, the value in the Undo Segment TX Slot event should be consistently at or near zero.

Option C:

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

Option D:

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

Question 39.

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

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

Answer: D

Explanation:

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

Option A:

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

Option B:

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

Option C:

Temporary tablespaces are not used to store extent allocation information.

Question 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.

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

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

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

Question 41.

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

- A. A group of table that each have more then 2 low cardinality columns.
- B. A data structure where a group of one or more tables have their own dedicated tablespaces.
- C. A group of one or more tables which resides in a tablespace that is striped across multiple disks.
- D. A group of one or more tables that share the same data blocks because they share common columns and are often used together in join queries.

Answer: D

Explanation:

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

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

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

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

Question 42.

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

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

Answer: B

Explanation:

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

Option A:

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

Option C:

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

Option D:

There is no compressed index type in Oracle.

Question 43.

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

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

Answer: C & E

Explanation:

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

Option A:

IOTs do not provide more concurrency..

Option B:

They will not be helpful for full table scans.

Option D:

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

Question 44.

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

- A. Materialized views cannot be partitioned, nor can they be defined on partitioned tables.
- B. Materialized views are often used in data warehouses to increase the speed of queries on very large datatables.
- C. Queries that benefit from the use of materialized views often involve joins between tables or aggregations such as SUM.
- D. A materialized view stores both the definition of a view and the rows resulting from the execution of the views.
- E. Materialized views can be used to replicate data, which was formerly achieved using the CREATE SNAPSHOT statement.

Answer: B, C, D & E

Explanation:

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

Option A:

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

Question 45.

Which statement is valid regarding index clusters?

- A. Index clusters can only be used for tables with low cardinality columns.
- B. Index clusters are generally well suited for tables that have many full table scans.
- C. Normal B-Tree indexes do not store null key values, whereas cluster indexes store null keys.
- D. A cluster index always takes up much more storage space than a normal index for the same set of key values..

Answer: C

Explanation:

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

Option 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.

Option B:

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

Option D:

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

Question 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.

Option A:

In statement 2 & 4 the index will be ignored.

Option B:

In statement 2 the index will be ignored.

Option C:

In statement 4 the index will be ignored.

Option D:

Statements 1 & 3 will use an index.

Option F:

Statements 1 & 3 will use an index.

Question 47.

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

- A. To avoid row chaining and migration.
- B. To speed up full table scans and fast full index scans.
- C. To improve performance when the index-organized table is clustered.
- D. To keep the B-Tree structure densely clustered to allow more rows per leaf block.

Answer: D

Explanation:

The main reason to use a row overflow for IOT is the need to keep the B-Tree structure densely clustered to allow more rows per leaf block. During IOT table creation you can specify the

OVERFLOW tablespace where the second half of the row data will be stored when the row's length exceeds the size set aside in PCTTHRESHOLD.

Option 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.

Option 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.

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

Question 47a.

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.

Option A:

This structure is not used to avoid row chaining and migration.

Option C:

IOTs are not good candidates to improve full table scans and fast full index scans. Row overflow is used for different reason.

Option D:

It does not improve performance when the index-organized table is clustered.

Question 48.

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

- A. Use the REFRESH COMPLETE option to truncate the materialized view and repopulate the view with data from the base tables in the query.
- B. Use the REFRESH FAST option to populate the materialized view data from the base tables that has changed since the last re-sync.
- C. Use the REFRESH FAST option to truncate the materialized view and populate the view with data from the base tables in the query.
- 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.

Option C:

The REFRESH FAST option does not truncate the materialized view.

Option D:

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

Option E:

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

Question 49.

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

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

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

Answer: A, C & E

Explanation:

Clusters, materialized views and index-organized tables will enhance database performance if they are used with appropriate types of applications. For example, Decision Support Systems (DSS) and data warehouse make heavy use of full table scans so the appropriate use of indexes and hash clusters are important. Index-organized tables can be also important tuning options for large DSS systems.

Option B:

Advanced Queue will not enhance database performance.

Option D:

Advanced Replication is not a data access method.

Option F:

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

Question 50.

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

- A. Bitmap indexes use less storage space.
- B. Bitmap indexes offer maximum concurrency.
- C. Bitmap indexes are easy to maintain when you issue DML statements.
- D. Bitmap segments are updated upon COMMIT, at the end of the transaction.

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

Answer: A & E

Explanation:

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

Option B:

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

Option 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.

Option D:

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

Question 51.

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

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

Answer: E

Explanation:

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

Option A:

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

Option B:

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

Option C:

There is no CHAIN segment exist for the IOT.

Option D:

There is no EXCESS segment exist for the IOT.

Question 52.

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

Which statement correctly describe one of the prerequisites?

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

Answer: A

Explanation:

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

Unlike the rule-based optimizer that has only its predefined guidelines to follow when executing a query, the *cost-based optimizer (CBO)* considers many different execution plans and then selects the one with the lowest execution plan. The QUERY_REWRITE_ENABLED init.ora parameter allows optimizer to dynamically rewrite queries to take advantage of materialized views when set to TRUE.

QUERY_REWRITE_INTEGRITY determines the degree to which the data consistency is to be adhered to when accessing materialized views. This will work only if OPTIMIZER_MODE is set to one of the cost-based optimizer modes.

Option B:

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

Option C:

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

Option D:

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

Option 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..

Question 53.

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

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

Answer: A & C

Explanation:

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

Option B:

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

Option 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.

Option E:

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

Question 54.

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

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

Answer: C.

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.

Option A:

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

Option 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.

Option D:

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

Question 55.

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

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

Answer: B

Explanation:

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

Option A:

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

Option 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).

Option 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.

Question 56.

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

- A. Oracle server allows query rewrites based on declared, but not enforced, relationships.
- B. Oracle server allows all updated materialized views and constraints with RELY flag to be used for the query rewrites.
- C. Query rewrites can occur 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.

Option 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.

Option 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.

Question 57.

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

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

Option B:

You cannot add or delete columns to partitions independently.

Option D:

You cannot add or delete constraints to partitions independently.

Option E:

You cannot change a column data type in each partition independently

Question 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.

Option A:

This information is stored in the corresponding locally managed tablespace tablespace, not in the memory.

Option B:

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

Option C:

Temporary tablespaces are not used to store extent allocation information.

Question 59.

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

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

Option A:

Frequent checkpoints cannot slow down a commit.

Option C:

They cannot cause the CKPT process to hang.

Question 60.

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

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

Answer: D

Explanation:

LOG_CHECKPOINT_INTERVAL initialization parameter set to 20 will achieve this goal: it shows number of OS blocks (20 * 0.5 K) when checkpoint will occur. LOG_CHECKPOINT_INTERVAL specifies the frequency of checkpoints in terms of the number of redo log file blocks that can exist between an incremental checkpoint and the last block written to the redo log. This number refers to physical operating system blocks, not database blocks.

Option 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.

Option B:

There is no LOG_CHECKPOINT_BYTES parameter in Oracle.

Option 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.

Option E:

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

Option 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.

Question 61.

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

```
RESOURCE_MANAGER_PLAN = day_oltp
```

What does 'day_oltp' specify?

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

Answer: A

Explanation:

RESOURCE_MANAGER_PLAN initialization parameter specifies a resource plan.

Option B:

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

Option C:

RESOURCE_MANAGER_PLAN initialization parameter does not specify a consumer group.

Option D:

Resource manager privilege cannot be set with this initialization parameter.

Question 62.

Why do these steps eliminate row migration?

Step 1: Run ANALYSE TABLE LIST CHAINED ROWS command

Step 2: Copy the rows to another table

Step 3: Delete the rows from the original table

Step 4: Insert the rows from step 2 back into the original table

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

Answer: A

Explanation:

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

Option B:

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

Option C:

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

Question 63.

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

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

Answer: B

Explanation:

Like Shared Servers, the performance of the Dispatcher process can also be monitored using a "busy ratio".

This ratio can be calculated using the V\$DISPATCHER dynamic performance view. Oracle recommends adding Dispatcher process if the Dispatcher busy rate consistently exceeds 50 percent. Another indication of Shared Server performance is the amount of time that User Process are waiting to have their requests accepted by the Dispatchers.

Option A:

Listener process waiting have nothing to do with dispatchers.

Option 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.

Option D:

Number of dispatchers is not related with dedicated connection process.

Question 64.

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

- A. A SECONDS_IN_WAIT value close to zero is ideal.
- B. A nonzero value in the SECONDS_IN_WAIT may indicate disk I/O contention on the redo log files.
- C. The SECONDS_IN_WAIT value of the log buffer space event indicated time spent waiting for space in the redo log buffer.
- D. A nonzero value in the SECONDS_IN_WAIT may be an indication the redo log buffers are too large and log switchers are not occurring fast enough.

Answer: A, 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.

Option B:

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

Question 65.

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

- A. DB_BLOCK_SIZE
- B. REDO_LOG_BUFFERS
- C. DB_BLOCK_BUFFERS
- D. DB_BLOCK_MAX_DIRT_TARGET
- E. DB_FILE_MULTIBLOCK_READ_COUNT

Answer: A & E

Explanation:

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

Option 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.

Option 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..

Option D:

There is no DB_BLOCK_MAX_DIRT_TARGET parameter in Oracle.

Question 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.

Option A:

IOTs do not provide more concurrency.

Option B:

They will not be helpful for full table scans.

Option D:

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

Question 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.

Option B:

This view contains information about file read/write statistics.

Option C:

This view contains information about current sessions.

Option D:

This view contains datafile information from the control file.

Question 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.

Option C:

There is no V\$SESSION_STATS view in Oracle.

Option D:

There is no V\$SESSION_STATUS view in Oracle.

Option E:

There is no V\$WAITS_PER_SESSION view in Oracle.

Question 69.

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

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

Answer: C

Explanation:

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

Option A:

This formula does not provide correct amount of memory. Parallelism can significantly improve a sorting time.

Option B:

SQL_AREA_SIZE should be multiplied on 2 to receive a correct formula.

Option D:

This formula incorrect: a total amount of memory for sorting is not related with query servers.

Option E:

This formula incorrect: a total amount of memory for sorting is not related with query servers.

Question 70.

How are deadlocks resolved within Oracle?

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

Answer: C

Explanation:

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

Option A:

The DBA should not terminate the waiting session: Oracle can handle deadlock situations automatically in most cases.

Option B:

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

Option D:

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