TROY TECHNOLOGIES USA

ORACLE 8 STUDY GUIDE

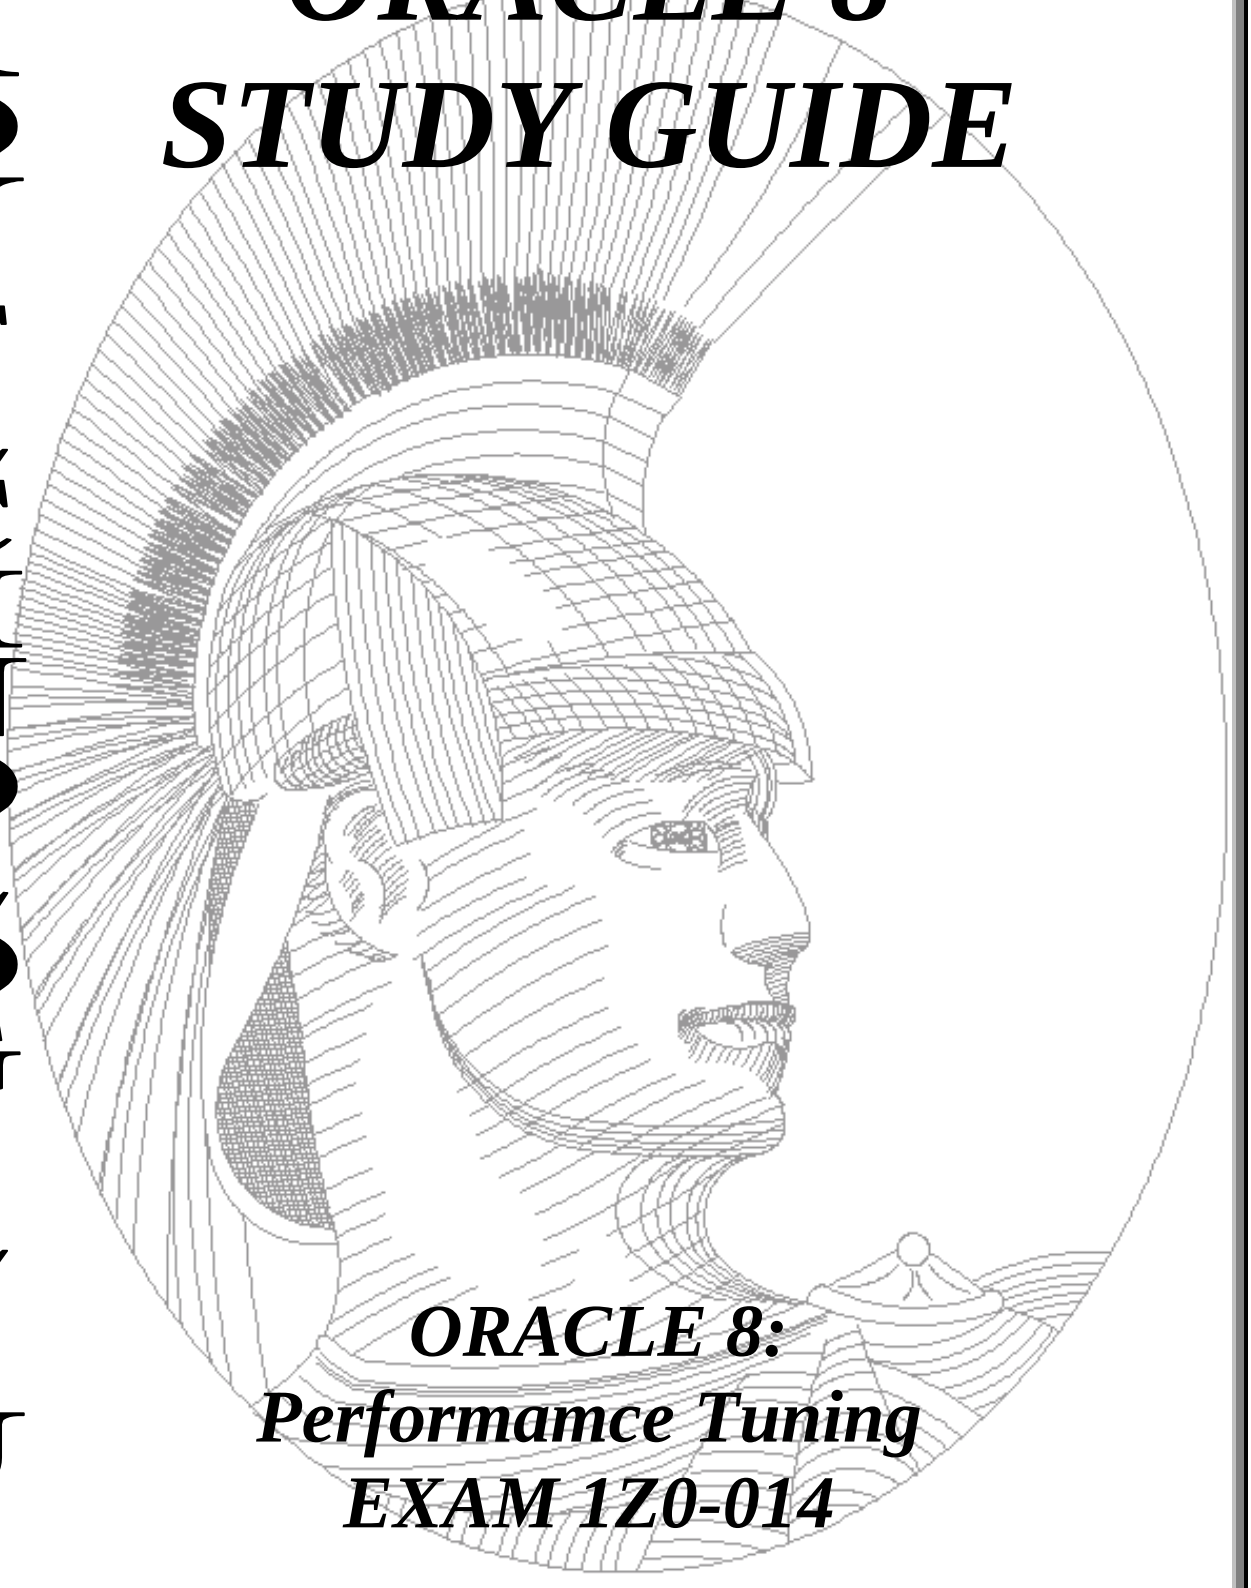

ORACLE 8: Performamce Tuning EXAM 1Z0-014

Edition 1

Congratulations!!

You have purchased a ***Troy Technologies USA*** Study Guide.

This study guide is a selection of questions and answers similar to the ones you will find on the official ORACLE 8: Perfromance Tuning 1Z0-014 exam. Study and memorize the following concepts, questions and answers for approximately 10 to 12 hours and you will be prepared to take the exams. We guarantee it!

Remember, average study time is 10 to 12 hours and then you are ready!!!

GOOD LUCK!

Guarantee

If you use this study guide correctly and still fail the exam, send your official score notice and mailing address to:

Troy Technologies USA
8200 Pat Booker Rd. #368
San Antonio, TX 78233

We will gladly refund the cost of this study guide. However, you will not need this guarantee if you follow the above instructions.

Table of Contents

Performance Tuning Key Concepts

Business Requirements and Tuning

Before beginning Oracle specific tuning you should check for bottlenecks on the operating system level. In particular, determine if there is excessive paging to the swap file and whether other applications are contending for system resources (such as RAM, CPU and disks) with Oracle.

Database Design

Many times you will not be able to participate at this level but when possible, being able to lay down a good design initially will help prevent a lot of performance nightmares from ever occurring.

Memory Tuning

Make sure the SGA is tuned correctly. The SGA is composed of three parts: 1) the data block buffer cache, 2) the redo log cache and 3) the shared pool. All three need to be tuned properly for optimal performance.

Disk I/O

Make sure your tablespaces are laid out correctly so that you don't have one type of tablespace adversely affecting another. An example of this is putting a rollback segment tablespace on the same disk as your data tablespace. Also verify you have storage parameters set optimally.

Internal Memory Structure Contention

Verify one internal process isn't waiting upon another for long periods of time. This is done by monitoring and adjusting latches correctly.

Oracle Alert, Trace Files, and Events

Trace files are Oracle log files which give the status of Oracle events. The main overall trace file which monitors the instance is called the alert.log. It is located in the directory specified by the BACKGROUND_DUMP_DEST initialization parameter. Background process (e.g., SMON, DBWR, etc.) trace files go here too.

Trace files (except for alert.log) usually end with a .trc extension. User process trace files go in the directory specified by the **USER_DUMP_DEST** initialization parameter. The initialization parameter **MAX_DUMP_FILE_SIZE** specifies the maximum size for user process trace files. In contrast, the alert.log trace file cannot have a specified size limit and must be truncated manually.

The V$ views are dynamic performance views based on the X$ tables. X$ tables are internal tables which hold information about the instance. Both are owned by SYS and populated at instance startup. V$ views can be seen by anyone with the "SELECT ANY TABLE" object privilege while X$ tables can only be viewed by SYS.

V$SYSTEM_EVENT, V$SESSION_EVENT, V$SESSION_WAIT - views which show information about processes waiting for resources.
V$SYSTEM_EVENT - shows total waits for events since the database started.
V$SESSION_EVENT - shows total waits for events since the database started grouped by session.
V$SESSION_WAIT - shows the most information about event waits per session. If the **WAIT_TIME** column of V$SESSION_WAIT has a value over 0 this indicates the session's last wait time. If the value is zero, the session is currently waiting. If the value is -1, the last wait was so small it wasn't measurable. If the value is -2, you need to enable **TIMED_STATISTICS** (see below) because Oracle isn't currently gather any time information.

Utilities and Dynamic Performance Views

Utilities

ULTBSTAT.SQL and UTLESTAT.SQL (B for begin and E for end) are scripts for collecting a variety of performance statistics. These scripts are located in $ORACLE_HOME/rdbms/admin.
Before running UTLBSTAT.SQL and UTLESTAT.SQL you should set the initialization parameter **TIMED_STATISTICS** equal to TRUE. This will give CPU time statistics. Alternatively you can enable and disable **TIMED_STATISTICS** on the fly by issuing: "**ALTER SYSTEM SET TIMED_STATISTICS=TRUE**" (or FALSE).

Report.txt contains the following sections:

1. Library cache statistics.
2. Numbers on a variety of statistics for the duration of the time UTLBSTAT.SQL was gathering data.
3. System wide wait events.
4. Average length of the dirty buffer write queue.
5. I/O statistics by datafile and tablespace.
6. Latch statistics.
7. Rollback segment statistics.
8. The current init<sid>.ora parameters.
9. Row (dictionary) cache statistics.

There is a set of GUI applications that come with the Oracle Enterprise Edition version of Oracle Enterprise Manager called the Oracle Performance Pack. These programs help monitor and optimize database performance:

Oracle Expert - assists in configuring and tuning your database (see below section for more detail).
Lock Manager - monitors locks.

Performance Manager - lets you monitor real time performance statistics and view them in various ways.
Tablespace Manager - lets you monitor segment storage in tablespaces. Also lets you defragment tablespaces (know as coalescing).
TopSessions - like the Unix 'top' command, lets you monitor the top resource intensive user sessions.
Oracle Trace - lets you monitor performance within user applications.

Performance Views

V$SYSSTAT is the main view for system performance on the instance level. You can also use V$SYSSTAT to monitor client-server traffic.

V$SESSION gives connection information for all user sessions.

Performance Ratios

- V$LIBRARYCACHE gives the pins/reload ratio for the library cache. The GETHITRATIO column should be .95 or higher. The RELOADS column should be equal to or less than one percent of the PINS column.
- The GETMISSES to GETS columns in V$ROWCACHE should have a ratio less than 15% percent.
- The buffer cache hit ratio should be 90% or higher.
- The hit ratio for a latch should be .98 or higher. Anything lower indicates latch contention.
- If in V$LATCH the MISSES to GETS ratio is higher than 1%, there is a problem with redo allocation latch contention.
- If in V$LATCH the IMMEDIATE_MISSES to IMMEDIATE_GETS ratio is higher than 1%, there is a problem with redo copy latch contention.
- Query V$WAITSTAT to see if there is rollback segment header block contention. A value over zero in the UNDO HEADER column (undo = rollback) indicates contention.
- Query V$SYSTAT to see whether sorts are occurring on disk or in memory. For OLTP systems 95% of sorts should be happening in memory.

Tuning Considerations For Different Applications

Know the difference between OLTP, DSS and hybrid systems. OLTP stands for online transaction processing system and is characterized by a high volume of inserts, updates and deletes (DML). DSS stands for decision support system. Is also commonly known as a data warehouse or OLAP: online analytical processing system. DSS is a read-only database consisting of massive amounts of data.

A *hybrid* system would be one that is both an OLTP database and a DSS database but at different times. An example of a hybrid database is a purchase order database that is used for querying during the daytime. But at night, new orders are inserted, current ones are modified, and fulfilled orders are deleted.

SQL Tuning

There are two optimizers built into Oracle: Rules Based and Cost Based. The Rules Based Optimizer (RBO) is a set of 15 rules Oracle uses to determine the fastest path of execution for a given SQL statement. The Cost Based Optimizer (CBO) determines all possible paths of execution and assigns a cost to each one. It chooses what it finds to be the least expensive path. You must be running ANALYZE on your tables and indexes to generate statistics to use the Cost Based Optimizer.

It is very important to update statistics on a regular basis to provide meaningful data for the CBO work with. Performance can suffer tremendously by using outdated statistics. You can set the CBO goal to either; 1) fastest overall throughput for a SQL statement or, 2) fastest initial response time for a SQL statement.

You can set the optimizer at the session level by issuing the command:

ALTER SESSION SET OPTIMIZER GOAL=*optimizer mode*

Use EXPLAIN PLAN FOR *sql statement* to see how the optimizer is executing your statement. This puts the execution plan in the PLAN_TABLE.

SQL Trace

SQL Trace is used to gather user session statistics. It generates a trace file which is generated in **USER_DUMP_DEST**. It contains the same data as AUTOTRACE with the following additional information:

- parse, fetch and execute counts
- cpu and elapsed time
- number of logical and physical reads
- number of rows processed
- library cache misses

SQL Trace can be turned on at the instance level by setting the initialization parameter SQL_TRACE=TRUE. SQL Trace can also be turned on for an individual session by issuing:

ALTER SESSION SET SQL_TRACE=TRUE;

You can also turn SQL Trace on in the current session by executing the procedure SYS.DBMS_SESSION.SET_SQL_TRACE(TRUE).

Finally, you can turn on SQL Trace for another session by executing the procedure:

SYS.DBMS_SESSION.SET_SQL_TRACE_IN_SESSION(*sid,serial#,*TRUE).

TKPROF is used to read the output of a trace file created by SQL Trace. The DMBS_APPLICATION_INFO package is created by the DBMSUTIL.SQL. It allows you to track resource usage and performance data for PL/SQL procedures. DBMSUTIL.SQL is called by the CATPROC.SQL script so it should already exist in your database.

Each lookup table has a *primary key* to *foreign key* relationship to the fact table (a child-parent relationship). However, in a star schema there is no relation between the various lookup tables among themselves. Hence, the "star" formation.

A *hash join* is a special type of equijoin that works well with joins between a large and small table. The initialization parameter **HASH_JOINED_ENABLED** controls whether hash joins are allowed. The default value is TRUE. This can also be changed at the session level via ALTER SESSION or at the statement level via the /*+ USE_HASH (*tablename*) */ hint.

Tuning the Shared Pool

The shared pool is composed of the library cache, row cache and if MTS is used, the User Global Area (UGA). The size of the shared pool is set by the initialization parameter **SHARED_POOL_SIZE** (in bytes).

The library cache contains shared parse information for SQL statements. This is the chief area to monitor in the shared pool.

V$LIBRARYCACHE - gives the pins to reload ratio for the library cache. The GETHITRATIO column should be .95 or higher. In V$LIBRARYCACHE, the value for RELOADS should be equal to or less than one percent of the value of PINS. If it is higher, you should increase the size of the shared pool.
V$SQLAREA - gives information about shared cursors and the first part of the SQL statements they contain.
V$DB_OBJECT_CACHE - gives information about PL/SQL packages and views. The SHAREABLE_MEM column specifies how much memory of the library cache an object is consuming.
V$SQLTEXT - gives the full SQL statements contained in shared cursors.
V$DB_OBJECT_CACHE - gives detailed information about cached tables, rows and PL/SQL objects.

The row (dictionary) cache stores date dictionary information.

V$ROWCACHE - view gives the hit/miss ratio for the dictionary cache. The GETMISSES to GETS columns in V$ROWCACHE should have a ratio of less than 15% percent. This ratio is expressed in report.txt as GET_MISS to GET_REQS.

You can pin packages, procedures, triggers and cursors into the shared pool using **DBMS_SHARED_POOL.KEEP( )**. The **DBMS_SHARED_POOL** package is created by running DBMSPOOL.SQL and PRVTPOOL.SQL.

The initialization parameter **SHARED_POOL_RESERVED_SIZE** specifies how much of the shared pool you want to set aside for the reserved list. This is an area of the library cache to store large objects in. Objects smaller that the value specified by the initialization parameter **SHARED_POOL_RESERVED_MIN_ALLOC** will not be allowed on the reserved list.

The User Global Area (UGA)is an additional third area of the shared pool when Oracle runs in Multithreaded Server (MTS) mode. When running in MTS mode, user information stored in the PGA in dedicated server mode is stored in the UGA instead. Hence, the UGA needs to be taken into consideration for overall sizing of the shared pool. Specifically you need to calculate the additional amount of memory required by the shared server sessions and open cursors.

Tuning the Buffer Cache

The size of the buffer cache is determined by: DB_BLOCK_SIZE * DB_BLOCK_BUFFERS. DB_BLOCK_SIZE cannot be changed without recreating the database. The buffer cache contains the dirty buffer write queue, which holds dirty block buffers (i.e., modified blocks) until they can be written out to disk. A least recently used (LRU) algorithm is used to decide which buffers to move out of the buffer cache so that new buffers can be read in.

DBWR is the background process that writes the dirty block buffers (located in the dirty buffer write queue) to disk. Server processes (S000..S999) read data blocks from data files into the buffer cache. Server processes check first to determine if the data blocks already exist in the buffer cache.

V$SYSSTAT - view that shows the *buffer cache hit ratio*. The buffer cache hit ratio should be 90% or higher.
V$RECENT_BUCKET - view that shows the effects of increasing the buffer cache by x amount of blocks. The value in the ROWNUM column shows the number of proposed block additions minus one (it starts with zero). The value in the COUNT column shows how many more buffer cache hits you will get by increasing the buffer cache by the proposed amount of blocks. To use V$RECENT_BUCKET, stop the database and set the **DB_BLOCK_LRU_EXTENDED_STATISTICS** initialization parameter to the maximum number of blocks you are considering adding. Next, start the database. After the instance has been running for a while under normal conditions query V$RECENT_BUCKET.
V$CURRENT_BUCKET - shows the effect of decreasing the buffer cache. The value in the ROWNUM column shows the number of proposed block subtractions. The value in COUNT column shows how many buffer cache hits you would still have after this decrease. To use V$CURRENT_BUCKET, stop the database and set the **DB_BLOCK_LRU_STATISTICS** initialization parameter to TRUE. Next, start the database. After the instance has been running for a while under normal conditions query V$CURRENT_BUCKET.

Gathering statistics for V$RECENT_BUCKET and V$CURRENT_BUCKET consumes CPU cycles and hence is a performance hit to the database. Therefore, you should disable these views by resetting their respective initialization parameters to the default value (i.e., 0 and FALSE) as soon as you are finished using them.

You should increase the buffer cache if:

1) The cache hit ratio is too low (below 90%).
2) There are increasing values for free buffer inspected (check V$SYSSTAT or report.txt).
3) There is contention for the cache buffer's chain latch (check V$LATCH).
4) There are buffer busy waits (check V$SYSTEM_EVENT or V$SESSION_WAIT).

On multi-CPU systems, **DB_BLOCK_LRU_LATCHES** specifies how many LRU latches (buffer chain latches) are available for the database. The default value is one half the number of CPUs and can be increased to double the number of CPUs. On single CPU systems you can only have one LRU latch.

A new feature in Oracle 8 is optional *multiple buffer pools* within the buffer cache. Using these buffer pools lets you explicitly define which objects you want to remain in the buffer cache and which you want to quickly unload. There are three buffer pools:

- Keep Buffer Pool: For objects you want to remain in the buffer cache as long as possible.
- Recycle Buffer Pool: For objects you want unloaded from the buffer cache as soon as the transaction that uses them is complete.
- Default Buffer Pool: The default buffer pool that DB blocks are stored in if not otherwise specified. Applications based on Oracle 7.3 and earlier will always use the default buffer pool. By default, the keep and recycle buffer pools are not enabled. To specify space for them, you must define the initialization parameters **BUFFER_POOL_KEEP** and **BUFFER_POOL_RECYCLE**. You must also specify how many buffer chain latches you want allocated for each pool. You must have at least 50 buffer blocks for every LRU latch you assign.

Syntax: BUFFER_POOL_KEEP(buffers:50000, lru_latches:5)

The size of the default buffer pool is not explicitly defined. Instead it is equal to the size of the entire buffer cache (value of **DB_BLOCK_BUFFERS**) minus the number of blocks allocated to the keep buffer pool and/or the recycle buffer pool.

The buffer blocks for the keep and recycle buffer pools are taken from the buffer cache. Likewise, their LRU latches are taken from the latches allocated for the entire buffer cache as specified by **DB_BLOCK_LRU_LATCHES**. Therefore, neither the number of buffer blocks nor the number of LRU latches allocated to the keep and recycle buffer pools can, taken together, equal or exceed the values allocated to the buffer cache as a whole. If you do over-allocate either DB blocks or LRU latches by accident, the database will not mount.

Use the CACHE hint (/*+ CACHE */) in your SQL statements to ensure that the selected table will be cached. You can also use the CACHE *clause* when creating a table to make sure it will always be cached.

Only cache small, frequently used tables. By default, large tables will fill the buffer cache (unless assigned to the keep or recycle buffer pool).

V$CACHE shows what objects are currently in the buffer cache. Run CATPARR.SQL to create this view. While CATPARR.SQL is for Parallel Server environments, V$CACHE is useful in single instance environments as well. The initialization parameter **CACHE_SIZE_THRESHOLD** specifies how many blocks per table will be cached. This is helpful to prevent buffer cache crowding.

Latch and Contention Issues

What locks are to tables, latches are to internal memory structures. Latches regulate access to instance resources by system and user processes. The V$LATCH view gives real time detailed latch information. The HIT_RATIO column gives the hit ratio for each latch. The SLEEPS column indicates how many times a process has had to wait to obtain a latch. Latch statistics are also listed in report.txt. The hit ratio for a latch should be .98 or higher. Anything lower indicates latch contention.

Tuning the Redo Log Buffer

The LOG_BUFFER initialization parameter specifies how big (in bytes) the redo log buffer is. Query V$SYSTEM_EVENT to see if there are waits for ‘log buffer space’. If so, increase the redo log buffer size. You can also obtain this information from V$SYSSTAT by looking at the number of 'redo buffer allocation retries' there are for a given user process.

Query V$SYSTEM_EVENT to determine if there are waits for ‘log file parallel write’. If so, this means there is I/O contention for redo log files. Stripe the redo log files across several disks to alleviate this problem.

You can forego creating redo log information during ALTER and CREATE statements as well as SQL*Loader direct-path loads by specifying the NOLOGGING clause. This greatly speeds up creation of large tables and indexes. NOLOGGING replaces the UNRECOVERABLE keyword used in Oracle 7.3.

There are two types of latches associated with the redo log buffer. The first is the *redo allocation latch*. The second type of latch is the *redo copy latch*. This latch exists only in multi-CPU systems.

The redo copy latch is used to break up the allocation and writing processes into two steps. The V$LATCH view gives statistics about redo allocation latch and redo copy latch contention. For redo allocation latch requests (WILLING_TO_WAIT type) the columns to monitor are GETS, MISSES and SLEEPS. For redo copy latch requests (IMMEDIATE type) the columns to monitor are IMMEDIATE_GETS and IMMEDIATE_MISSES.

If the MISSES to GETS ratio is higher than 1%, there is a problem with redo allocation latch contention. If the IMMEDIATE_MISSES to IMMEDIATE_GETS ratio is higher than 1%, there is a problem with redo copy latch contention.

Database Configuration and I/O Issues

Set the **LOG_CHECKPOINTS_TO_ALERT** initialization parameter to TRUE so that checkpoint beginning and ending times are logged to the alert.log file.

The **DB_BLOCK_CHECKPOINT_BATCH** specifies the maximum number of blocks that a DBWR process can write in a single batch during a checkpoint. Increasing this value can speed up checkpoint times, but making it too large can give poor response times as well.

The initialization parameters **DISK_ASYNCH_IO** and **TAPE_ASYNCH_IO** specify whether the operating system supports asynchronous I/O for hard drives and tape drives respectively (most do). The default value is TRUE.

If your operating system doesn't support asynchronous disk and/or tape I/O, you can enable them on the Oracle level by specifying the following initialization parameters:

- **ARCH_IO_SLAVES**
- **LGWR_IO_SLAVES**
- **DBWR_IO_SLAVES**
- **BACKUP_DISK_IO_SLAVES**
- **BACKUP_TAPE_IO_SLAVES**

The SYSTEM tablespace should only contain data dictionary objects, PL/SQL packages, triggers and the initial rollback segment. Your database should have a minimum of six tablespaces:

SYSTEM - for the data dictionary and initial rollback segment.
DATA - for tables.
INDEX - for indexes.
TEMP - for sorting.
RBS - for rollback segments.
USER - for user created tables.

V$FILESTAT tells the number of reads and writes done to each data file. Use this view and report.txt to determine if I/O is evenly distributed or not.

Avoid heavy I/O on the SYSTEM tablespace. Heavy I/O may occur by not specifying the default and temporary tablespace when creating users. If so, users object and query sorts will be written and read from SYSTEM.

If the row cache isn’t big enough to allow the entire data dictionary to exist in memory, additional I/O will occur in the SYSTEM tablespace due to data dictionary reads.

When using hardware or software striping, make sure the stripe block size is a multiple of the **DB_FILE_MULTIBLOCK_READ_COUNT** initialization parameter. Manually stripe tablespaces only if hardware or the operating system doesn't support it. You should only consider it if you are in an OLAP environment where a lot of full table scans are happening and you are using Parallel Query.

You can manually stripe by spreading the tablespace over two data files, one on a different drive. Set MINEXTENTS greater than 1 and then create the table so that each extent fills up more than half of each data file. You can also allocate extents to data files explicitly.

Using Oracle Blocks Efficiently

PCTUSED is relevant for deletes. If you have a lot of inserts and deletes on a table, set PCTUSED high. On DSS systems PCTUSED is irrelevant. PCTFREE and PCTUSED taken together should be less than 100.

Migrated and chained rows cause excessive disk I/O. So does inefficient use of block space which causes more blocks to be read than is necessary. The goal is to determine how your tables are being used and set the block storage parameters optimally, neither overcrowding nor wasting space.

The CHAIN_CNT column of **DBA_TABLES** tells you how many chained or migrated rows a table has. The CHAINED_ROWS table gives detailed information about each chained or migrated rows in a table. This table is created by the UTLCHAIN.SQL script. Use the LIST CHAINED ROWS clause of the ANALYZE command to populate the CHAINED_ROWS table.

DB_BLOCK_SIZE should be a multiple of the OS block size. Make it as large as possible.
Set **DB_FILE_MULTIBLOCK_READ_COUNT** to a multiple of **DB_BLOCK_SIZE**. This determines how many blocks are read at once in a full table scan. This is very important for DSS systems.

The *highwater* mark for a table is increased 5 blocks at a time. Full table scans read up through the highwater mark. DELETE does not reduce the highwater mark count on a table , but TRUNCATE does. So does ALTER TABLE *tablename* DEALLOCATE UNUSED.

Optimize Sort Operations

The following clauses cause sorting to occur: ORDER BY, GROUP BY, DISTINCT, UNION, INTERSECT and MINUS. Oracle will sort in memory if the sort can fit in the value specified by the initialization parameter **SORT_AREA_SIZE** (in bytes). If not, Oracle will break the sort into multiple *sort runs*. MAXEXTENTS is not a valid storage parameter for temporary tablespaces.

When a process needs to sort on disk it looks in the sort extent pool (SEP) in the SGA to get a free extent. V$SORT_SEGMENT gives detailed information regarding sort extents. The EXTENT_HITS column tells how many times a free segment was found in the sort extent pool. Query V$SYSTAT to see whether sorts are occurring on disk or in memory. At least 95% of sorts should be happening in memory for OLTP systems.

The initialization parameters **SORT_WRITE_BUFFERS** and **SORT_WRITE_BUFFER_SIZE** specify how much memory to allocate for direct write sorts. The **SORT_WRITE_BUFFERS** should be between 2 and 8. **SORT_WRITE_BUFFER_SIZE** should be between 32K and 64K.

Rollback Segment Tuning

Rollback segments are used for the three Rs: transaction rollback, instance recovery and query read consistency. There is a *Rule of 4* for sizing the number of rollback segments: # of concurrent transactions divided by 4. If less than 4+4, round up to a multiple of 4. No more than 50 total.

If you get an "ORA-01555**:** snapshot too old (rollback segment too small**)**" error, this means a rollback segment is too small for a given transaction to complete and that it has overwritten itself. You should increase the number of rollback segments or assign the transaction a large rollback segment.

Use "SET TRANSACTION USE ROLLBACK SEGMENT *rollback segment name*" to assign a specific rollback segment to a transaction. This should be done for long running batch jobs. MINEXTENTS for rollback segments should be 20. INITIAL for rollback segments should be set to the power of 2. OPTIMAL should be set high enough to accommodate normal database activity.

DELETEs use the most rollback space. Next UPDATEs, then INSERTs.

V$SESSION and V$TRANSACTION show how much rollback information a transaction is creating.

The rollback transaction table is in the rollback segment header. Query V$WAITSTAT to see if there is rollback segment header block contention. A value over zero in the UNDO HEADER column (undo = rollback) indicates contention.

Monitoring and Detecting Lock Contention

Enqueues are what Oracle uses to manage locks. DML locks occur at the row level. The V$LOCK view gives detailed lock information.

A query does not hold a lock unless the SELECT FOR UPDATE statement is issued. The SELECT FOR UPDATE statements give the user a Row Shared (RS) lock. This gives a shared table lock and an exclusive row lock. A Share lock on a table prevents DML operations happening to it. A Share Row Exclusive lock prevents DML and SELECT...FOR UPDATE on a table.

Query V$SESSION to see which row is causing lock contention. DDL locks are not usually in contention because they are held only briefly.

The three types of locks mainly in use by applications are: TM (row exclusive table lock), TX (row exclusive lock) and UL (PL/SQL user locks). High level locking (i.e., locking a whole table instead of just a row) in application code can cause the majority of lock contention in a database. This may be because the application was coded for another rdbms such as Sybase or SQL Server.

To kill a user session which is holding a lock, query V$SESSION and get the SERIAL# and SID (session ID). Next, issue ALTER SYSTEM KILL SESSION *serial #, sid*.

UTLLOCKT.SQL shows a graph of which locks are being held. CATBLOCK.SQL must be run first to set up the views for this.

When a deadlock is detected, Oracle will roll back the statement that caused the deadlock and issue an ORA-00060 error message. Oracle does not rollback the entire transaction, only the statement that detected the deadlock. You will have to manually rollback the rest of the transaction if necessary. When a deadlock occurs the rowid of the locking row is recorded in a trace file located in **USER_DUMP_DEST**.

Tuning with Oracle Expert

Oracle Expert uses the concept of a *tuning session* to focus in on a given portion of the database you want to examine. All data is collected and stored on the tuning session level. These sessions are saved to disk with an .XDL extension. There are three scopes for tuning sessions:

1) *Application level* - examines SQL statements and access methods (index usage).
2) *Instance level* - examines SGA, operating system and I/O parameters.
3) *Structure level* - examines storage parameters and placement of tablespaces.

Oracle Expert has seven different classes of input data:

1) *Database* - users, tablespaces, public synonyms.
2) *Instance* - data gathered from the X$ tables and V$ views.
3) *Schema* - tables, indexes, sequences, views, etc.
4) *Environment* - physical server attributes (RAM, CPU, disks).
5) *Workload* - what normal transactions are performed on the database. Can be gathered by Oracle Trace.
6) *Rules* - guidelines to tell Oracle how to treat collected data. Change rules if you don't like a given recommendation.
7) *Control parameters* - specifies whether the database is an OLTP, DSS, or hybrid, so Oracle Expert can set goals and use specific features accordingly.

Oracle Expert generates two different types of output:

1) *Reports* - includes a) Analysis, b) Session Data and c) Recommendation Summary.
2) *Implementation Files* - includes proposed changes to the init<sid>.ora file and SQL scripts to put into effect recommended changes (e.g. index creation).

Practice Questions

1. What is the purpose of the DB_BLOCK_LRU_EXTENDED_STATISTICS parameter?

A: Allows evaluation of the effects of adding buffers to the database buffer cache.

2. When creating a recycle buffer pool what should you use as a sizing guideline?

A: It should be large enough to retain blocks for the duration of the transaction.

3. Which indicator is the most useful measure of database buffer cache performance?

A: Cache hit ratio

4. Which relationship is used in computing the cache hit ratio?

A: 1-(physical reads/ (db block gets + consistent gets))

5. The V$CACHEVIEW is most useful when using which of the following?

A: Parallel Server

6. Your database utilizes small lookup tables, each of them is being accessed by many users. What can you do to improve the performance of full table scans?

A: Use the cache clause

7. Which two values are needed to compute the size, in bytes, of the database buffer cache? (CHOOSE TWO)

A: DB_BLOCK_BUFFERS
DB_BLOCK_SIZE

8. You are the DBA for an OLTP system. The amount of activity has increased considerably and you need to increase the size of the REDO LOG BUFFER. Which initialization parameter would you alter in order to increase the size of REDO LOG BUFFER?

A: LOG_BUFFER

9. You are performing a direct load to the TEMP_DATATABLE using SQL*Loader and you issue this command prior to the task:

ALTER TABLE temp_date NOLOGGING;

What has been accomplished?

A: Redo log entries will not be generated for the load to the temp_data table.

10. Which view can you query to determine if there are Waits occurring for redo log buffer space?

A: V$SESSION_WAIT

11. You are testing a new database to see how it will perform with a larger block size. Which situation would best allow you to perform this test?

A: You need a test database that you can easily duplicate changing the DB_BLOCK_SIZE parameter with new variations.

12. You have an application and you do not ever want users to experience delays while an extent is being allocated dynamically in several tables. What can you do?

A: Monitor the database for these segments and extend them manually.

13. Many segments are being extended dynamically. Which storage parameter should you alter to reduce the frequency of extent allocation?

A: MAXEXTENTS

14. Examine these commands:

```
SQL>ANALYZE INDEX acct-no-idx VALIDATE STRUCTURE;
Index Analyzed.
SQL>SELECT(DEL-LF-ROWS-LEN/LF-ROWS-LEN)*100
        As index-usage
                FROM index-stats.
                WHERE index-name= 'ACCT-NO-IDX'*;
INDEX USAGE
----------------------------------
24
```

What do you do to improve the condition of the index?

A: Rebuild the index.

15. If the emp table is updated frequently which formula should you use to determine the new value for the PCTFREE?

*A: PCTFREE = 100*a/(a-b) where a is the row size after expansion and b is the initial row size.*

16. You have been troubleshooting the PROD database and after issuing the ANALYZE command against several heavily used tables determined that row migration is heavy. What effect will these tables containing a large amount of migrated rows have on your database?

A: Increase I/O

17. Which two values would you retrieve from V$SESSION to issue an alter system kill session statement? (CHOOSE TWO)

A: SID
SERIAL#

18. Where are the trace files for dead lock situations recorded?

A: USER_DUMP_DEST

19. Which actions could you take to eliminate implicit share locking on a child table where delete statements are issued against the parent table?

A: Create an index on the foreign key column in the child table.

20. Users are testing a new application and you want to gather statistics for the database. Which script is provided by Oracle to begin gathering database statistics?

A: utlbstat.sql

21. Which oracle diagnostic tool set includes a coalescing function that allows you to join adjacent free blocks?

A: Oracle Tablespace Manager

22. Which parameter would you set first to obtain more useful information when running the utlbstat/utlestat utility?

A: TIMED_STATISTICS=TRUE

23. Performance has degraded and paging and swapping are occurring on the system since you increased the size of the SGA. Which condition might cause this?

A: Buffer cache is too big.

24. You have turned on the SGA to improve performance. In addition, you have set the PRE_PAGE_SGAINITIALIZATION parameter to true. What is the result of this setting?

A: All SGA pages are accessed and brought into the memory during instance startup.

25. Which SQL operation will not cause a sort?

A: UNION ALL

26. What happens to the sort memory area when a sort has finished its work?

A: The area shrinks to size specified by SORT_AREA_RETAINED_SIZE.

27. The new DBA adds a user to the database you later discover that he did not assign the temp tablespace as the allocation for this table sort area. To which tablespace does the sort area default?

A: SYSTEM

28. Which two storage parameters should be set as integer multiples of SORT_AREA_SIZE when creating the temporary tablespace? (CHOOSE TWO)

A: INITIAL
NEXT

29. Which view can be queried to find the total rows sorted?

A: V$SYSSTAT

30. You are monitoring the instance for redo latch contention and query V$LATCH. It shows a latch miss ratio of 6%. What does this value tell you about the condition of the redo allocation latches?

A: Contention is high.

31. You suspect redo latch contention in the PROD instance and query V$LATCH and V$LATCHNAME to obtain statistics. What should the ideal percentage of latch contention time be?

A: Less than 1%

32. Processes are in contention for access to the free list on the orders table. Which action should you take to eliminate or reduce this contention?

A: Drop and recreate the orders table, increase the number of free lists.

33. You are monitoring rollback segments usage for the PROD database to determine if you need to add additional rollback segments. Frequent queries of V$WAITSTAT during normal processing hours indicates an average value of 115 for the undo header column. What does the average value of this column indicate to you?

A: Contention exists for rollback segment header blocks in the transaction table.

34. Your supervisor has given you the responsibility for making sure your system rollback segments are correctly tuned. Which three goals should you set for this task? (CHOOSE THREE)

A: Rollback segments will not extend during normal running.
Transactions will never wait for access to rollback segments.
Readers will always see the read consistent images they need.

35. If your users are running long queries that access frequently changing data which aspect of rollback segment tuning should you address?

A: Make sure rollback segments do not wrap around and prevent the construction of a read consistent view.

36. What is the recommended default number of transactions per rollback segment in an OLTP environment?

A: 4

37. When sizing rollback segments, what is the recommended setting for the MAX EXTENTS parameter?

A: Any power of 2

38. How can dynamic space allocation be avoided in an OLTP system?

A: Explicitly preallocate space to tables, indexes and clusters.

39. Several B_TREE indexes exist and you believe performance for these indexes has degraded. After issuing the ANALYZE INDEX<I_NAME>COMPUTER STATISTICS, you determine that BLEVEL has a value of six. At what value should you keep BLEVEL for B_TREE indexes?

A: Less than 4.

40. What should you do if you need to tune a SQL statement that includes a tree walk?

A: Index the columns in the START WITH and CONNECT BY clauses.

41. Payroll application users are complaining about the length of time it takes to run the earnings calculation program, you have enabled tracing for this process and the trace is now complete. What would you use to format information about the execution of the SQL statements for this process?

A: TKPROF

42. When setting up a data warehouse to use star queries, which two objects would need to be structured properly for the star queries to be recognized by the cost-based optimizer? (CHOOSE TWO)

A: Fact tables
Lookup tables

43. You have just implemented the new financials application and want to track usage of various modules. What can you use to track the usage of modules in the application?

A: DBMS_APPLICATION_INFO

44. You have chosen to checkpoint only at log switches. To ensure this, which initialization parameter should you set to a number of operating system blocks larger than your redo log files?

A. LOG_CHECKPOINT_INTERVAL.

45. You have the database that has DB_BLOCK_SIZE set for 2KB. However you are planing to migrate the database to a new system that can accommodate a block size of 8KB. What can you do to take advantage of this enlarge block size?

A: Recreate the database on the new system with DB_BLOCK_SIZE set to 8KB.

46. You query INDEX_STAT because you suspect there is a problem with the ACCT_NO_IDX index. What information is returned?

A: Information from the last analyze index. . . validate structure command issued.

47. The PROD database was created with a data file named data01.dbf and you have since created data02.dbf on another device. How can you take advantage of manual stripping for these data files?

A: Choose the data file with the most available space.

48. Which data must be included for Oracle Expert to perform access methods tuning analysis?

A: Schema class data

49. What information is included in the analysis report generated by Oracle Expert for a tuning session?

A: Listing of all attributes for the instance for which you want tuning recommendations to be generated.

50. In preparing Oracle Expert with needed data what can only be entered manually or imported from .xdl file?

A: Workload class data

51. Some application processes are running slowly and you decide to enable SQL tracing. At which two levels can you choose to enable SQL tracing? (CHOOSE TWO)

A: SESSION
INSTANCE

52. When setting up the initialization file for an instance, which parameter would you set to control the location of a user trace file?

A: USER_DUMP_DEST

53. You are setting initialization parameters for the PROD instance and you need to set the location of the alert log file. Which initialization parameters should you set?

A: BACKGROUND_DUMP_DEST

54. You have been asked to develop a tuning plan for an application that has been in production for a year. What should be your most important tuning goal?

A: Tune all elements necessary for users to see noticeable performance improvements.

55. Your company is planning to design and implement a custom payroll application. At which time during this process is it most important for performance to be addressed?

A: Application development

56. Which part of a shared pool contains information about object privileges?

A: Dictionary cache

57. You are running MTS and need to properly size the UGA. Which information must be computed?

A: Required space for all user global areas to be placed into the memory.

58. Which parameter will you set to accommodate large objects in the memory cache?

A: SHARED_POOL_RESERVED_SIZE

59. You have determined that the payroll application processes are fragmenting the shared pool. What should you pin in the shared pool using the DBMS_SHARED_POOL packets to reduce fragmentation?

A: Large PL/SQL objects

60. What is an ideal value of the GETHITRATIO column in the V$LIBRARY CACHE?

A: Greater than 0.9

61. When are the packages loaded into library cache?

A: As soon as any part of them is accessed.

62. You are changing the storage parameter for the EMP table by setting PCTUSED 250. When will the change to this storage parameter be issued?

A: Immediately

63 The emp table is being heavily used and you want to obtain information about data storage. Which command can you issue to gather detailed storage information for this table?

A: ANALYZE

64. You run the utlbstat/utlestat twice to gather performance statistics for two intervals. But the output file contains only statistics for the last interval. What must you do to retain statistics for both intervals?

A: Rename the output file after running the utlbstat/utlestat utility the first time.

65. Which parameter should be increased if the RELOADS/PINS ratio is higher than 1%?

A: SHARED_POOL_SIZE.

66. You have pinned a large PL/SQL package in the shared pool and another DBA issues the alter system flush SHARED_POOL command. What happened to the PL/SQL package?

A: It is still in the shared pool.

67. In changing your database to use MTS, which initialization parameter may need to have its value increased?

A: SHARED_POOL_SIZE.

68. Examine this command:

SQL>execute DBMS_SHARED_POOL.KEEP('My_package');

What occurs as a result of issuing this command?

A: My_package is pinned in the shared pool.

69 If you set your database to use MTS where will the user session and cursor state information be stored?

A: UGA

70. Which query will return the percentage of parse cause that finds a cursor to share?

A: SELECT gethitratio
FROM v$librarycache
WHERE namespace= 'SQL AREA';

71. What are two reasons you should keep rollback segments in tablespaces dedicated to rollback segments? (CHOOSE TWO)

A: You can't take a tablespace off line unless it contains an active rollback segment.
The rollback segment may fragment.

72. You have a large recurring transaction but have only one rollback segment that is large enough to handle the transaction. How can you make sure this transaction uses the large rollback segment?

A: Issue the SET TRANSACTION USE ROLLBACK SEGMENT command.

73. Transactions are usually frequent and small but the processes for the financial applications generate some large transactions. How should rollback segments be assigned for this type of transaction?

A: Assign several small rollback segments to the transaction.

74. What should be the sum of Waits to the sum of Gets for rollback segment usage?

A: Less than 5%.

75. What is the minimum number of extends required in a rollback segment?

A: 2

76. Examine this query and results:

```
SQL>SELECT sid, event, seconds_in_wait, state
        FROM v$session_wait
        WHERE event like 'log buffer%';

SID   EVENT              SECONDS_IN_WAIT     STATE
5     log bufferspace    110                 WAITING
```

What might the results indicate?

A: redo log buffer is too small.

77. You created the test database recently and now you think you may not have made the redo log buffer large enough. Which query would you use to see if Waits are occurring for the event log buffer space?

A: V$ SYSTEM_EVENT.

78. You are performing a direct path to load to the emp_data table in the PROD database which is set to the NOARCHIVE mode. How will a redo be generated for this load?

A: Redo will be generated only for the data dictionary changes.

79. Why should you avoid using the DISTINCT keyword in statements that select a large number of rows?

A: The resulting sort can mean an increase in I/O and processing overhead.

80. What is best value of PCTINCREASE for the temporary tablespace?

A: 0

81. Which parameter defines whether sort writes will use the buffer cache?

A: SORT_DIRECT_WRITES.

82. What would you do to help ensure that sorting is done in memory?

A: Increase the value of the SORT_AREA_SIZE parameter.

83. You issue the statement CREATE INDEX. . . NO SORT. What must be done to the table data you wish to index?

A: The data will have to be loaded in sorted order.

84. Which three things should you check to see the files in a database with the most I/O activity? (CHOOSE THREE)

A: V$FILESTAT
FILE I/O using Oracle Performance Manager
File statistics in report.txt

85. You have determined that row migration is occurring for several tables in the PROD database. Which storage parameter can you set to prevent further migration for the involved tables and any other tables where you do not want this to occur?

A: PCTFREE

86. In an attempt to resolve redo latch contention, you increase LOG_SIMULTANEOUS_COPIES from 2 to 4. What have you accomplished?

A: Increase the number of redo copy latches.

87. You suspect there is I/O contentions with the log files. Which view would you query for confirmation?

A: V$LOGFILE

88. What does Oracle 8 use to regulate access to internal memory structure such as the library cache?

A: Latches

89. In your efforts to tune the whole system you want to provide the system administrator with statistics to determine where to move data files and how to balance I/O. With which view can you gather I/O statistics?

A: V$FILESTAT

90. Which three system statistics are used to compute the database buffer cache hit ratio? (CHOOSE THREE)

A: DB block gets

Consistent gets
Physical reads

91. You have created keep, recycle and default buffer pools for the PROD instance. However when you start up the instance an error occurs and the database is not mounted. What would cause this?

A: Total allocation to the buffer pools has exceeded the DB_BLOCK_BUFFER and/or DB_BLOCK_LRU_LATCHES parameters.

92. Which is the best way to minimize the number of blocks visited to retrieve data in your existing database?

A: Tune the SQL statement

93. Users are reporting slow response time on your system. You think that the problem can be solved by increasing the size of the database buffer cache. You want to make sure your solution is correct. Which two steps are necessary to evaluate the impact of adding buffers to the database buffer cache? (CHOOSE TWO)

A: Set the parameter DB_BLOCK_LRU_EXTENDED_STATISTICS to the number of buffers you want to add.
Query the V$RECENT_BUCKET view.

94. How might you prevent blocks from going to the least recently used end of the LRU list?

A: Create or alter tables using the cache clause.

95. In your efforts to improve performance you altered several frequently accessed small tables with the cache clause. Which initialization parameter can you set to prevent over crowding in the buffer cache?

A: CACHE_SIZE_THRESHOLD.

96. Which value is affected by changing the DB_FILE_MULTIBLOCK_READ_COUNT parameter?

A: The number of blocks read at one time during a full table scan.

97. You have enabled SQL Tracing at the session level. What will this trace file contain?

A: Statistics for SQL statements for the session.

98. You are the DBA for a busy OLTP reservation system. Which two parameters can you set to increase the overall numbers of locks available for this instance? (CHOOSE TWO)

A: DML_LOCKS
ENQUEUE_RESOURCES

99. The Oracle server automatically detects and dissolves deadlocks by always rolling back, what?

A: The transaction that detected the deadlock.

100. Your decisions support system is performing poorly when data is queried. After analyzing the database to improve query performance you determine that an index could be added on a column that has few distinct values. Which index type should you create?

A: BIT MAPPED

101. Performance for queries in your decision support system has degraded and you decide to create indexes for the relevant tables. The tables are read only and the queries usually involve multiple predicates on a low cardinality column. Which index type should you create?

A: BITMAPPED

102. What can you use in SQL*Plus to find out how a statement would be executed if you ran it at that moment?

A: EXPLAIN

103. You have been monitoring the account receivable application and notice some queries are taking along time to complete. Using SQL trace and TKPROF you find out the full table actions are performed. Which low level action can you take to possibly improve performance for this process and decrease the load?

A: Add appropriate indexes.

104. Payroll application users are complaining about the length of time it takes to run the earning calculations process. What would you use to gather detailed information for this process while it is running?

A: DBMS_APPLICATION_INFO.

105. You query V$SYSTEM_EVENT and a value of 15 is displayed for the TOTAL_WAITS column. For which period of time would this value has been determined?

A: Since instance start up

106. What type of information does the alert log file contain?

A: Messages and errors.

107. 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

108. Which two statements can cause dynamic extents? (CHOOSE TWO)

A: Insert
Update

Index