Database Performance Utility

Author: G S Chapman

Date: 15th March 2005

Version: 1.1

Location of Document:

DOCUMENT HISTORY

Version	Date	Changed By:	Remarks
1.0	01/02/05	G S Chapman	Initial Version
1.1	15/02/05	G S Chapman	Updated with extra testing instructions

DOCUMENT DISTRIBUTION

Copy No	Name	Role	Organisation

DOCUMENT REFERENCES

Document Name	Originator	Part Number	Version	Date

TABLE OF CONTENTS

TABLES

Appendices

PURPOSE OF DOCUMENT

This document describes a generic database performance package and associated packages to be employed upon Oracle databases. The current release also introduced a dependency upon Oracle 10g. See section E.1 below for more details.

Overview

1 Introduction

There has always been a requirement to perform database monitoring to ensure that a database is running to its best possible efficiency. Traditionally before Oracle 8.1.6 (8i) the only method of performing this task was to employ 'user written scripts' that would be run periodically by the Database Administrator (DBA).

The 'user written' scripts varied is shape and form and in their general usefulness and there was no general standard that could be recognised from site to site or even in some cases database to database.

To address this requirement a number of scripts were developed, and used to enable a small group of DBA's to have familiarity between databases where a number of different clients were concerned. This enabled easier and quicker problem determination.

With the release of Oracle 8.1.6 an Oracle supplied packages called 'statspack' was available, which became a part of the official database release with Oracle 8.1.7. Further enhancements were made to 'statspack' and with Oracle 10.1 (10g) this has been enhanced with the introduction of ADDM.

Oracle Database 10*g* introduced an integrated set of self-managing capabilities to simplify administration, increase efficiency and lower the costs associated with systems management, whatever the workload.

The components of the performance diagnosis and monitoring technology are built into the database server and externalised through Oracle Enterprise Manager (EM). The major components are the Automatic Workload Repository (AWR), Automatic Database Diagnostic Monitor (ADDM), and Oracle Enterprise Manager (EM). Underlying all of the components is the code instrumentation in the Oracle database code that generates the wealth of diagnostic statistics available from the Oracle database.

2 Generic Databases

Very often, as is common on a number of database sites, the database administrator is involved in a number of activities not related to the database and consequently a lot of the time the database is left to its own devices.

The DBA only really gets involved when a problem is reported, such as poor response etc. The wealth of tools are available with Oracle 10g accessed through the Enterprise Manager (EM) are not used to their best advantage and an easier method of database monitoring, that will alert concerned parties in the event of possible problems is required.

3 STATS Package

To address the identified needs the 'STATS' package was written and has been enhanced with new functionality to make it generically useful upon many possible databases.

The package has the ability to be configured to work with any version of the Oracle database, and reports can be created in a number of ways:

- written to flat files within the operating system.
- emailed directly to recipients (assuming the site security allows such functionality).

- alerts sent via SMS (paging) (requires access to an SMS gateway.) Currently unimplemented.
- A possible further enhancement might be to implement an 'ftp' file transfer.

Newer reporting features will of course depend upon the version of the database upon which the package is installed. Thus the ADDM and AWR reports will not be available with pre Oracle 10g database versions.

The package itself calls other generic packages. In particular these are the EMAIL and the LOGGING packages. These additional packages are very open in nature to enable their easy reuse in other packages as required. They are documented and described below to ensure completeness.

Performance Tools

4 Performance report pre Oracle 8.1.6

Before Oracle 8.1.6 there were no Oracle supplied tools for measuring database performance. Instead each DBA would have their own 'written' set of scripts which, they would use to monitor the databases for which they were responsible.

The set that were used for the routines within the 'STATS' package were built up over a number of years and have been found extremely useful. They have been used for many years, by Oracle DBA's. The scripts would save certain pieces of information, which were mainly the database buffer cache, library cache, etc type of statistics to enable a trend analysis to be performed. The keeping of these statistics was however optional and it has always been possible to just have a 'one-off' run.

The output from the 'packaged' version of the original scripts has deliberately written to match that obtained from the original scripts. The packaged version has been enhanced with additional functionality as time and circumstances permit but the original scripts have not been updated for some time.

The schema creation script needs to be run initially in the desired databases prior to the running of the scripts as part of the installation.

Note that there is no need to connect to the database as the package owner, and the user can use their usual DBA-privileged account to run the scripts.

5 Oracle Statspack package

Statspack is a performance diagnosis tool, available with Oracle8i Release 8.1.6. Statspack can be considered BSTAT/ESTAT's successor, incorporating many new features, such as:

- Storing performance data in permanent Oracle tables which can be used for analysis at any time in the future
- Capturing high-resource SQL statements
- An easy to use report, with many useful ratios pre-computed
- The ability to be run on all nodes of an OPS system

Statspack is a diagnosis tool for instance-wide performance problems; it also supports application tuning activities by providing data which identifies high-load SQL statements. Statspack can be used both proactively to monitor the changing load on a system, and also reactively to investigate a performance problem.

The Statspack output is similar to that generated by AWR in Oracle 10g.

6 Oracle 10g Database Diagnostics

The Automatic Database Diagnostic Monitor (ADDM) built into the Oracle Database 10g provides the following benefits:

- Automatic performance diagnostic report every 60 minutes
- Problem diagnosis based on decades of tuning expertise
- Time-based quantification of problem impacts and recommendation benefits
- Identification of root cause, not symptoms

- Greatly reduced need to replay workload for detailed analysis due to completeness of the data held in the Automatic Workload Repository (AWR).

6.1 Database Statistics

With each new release of the database more performance statistics are added that allow us to diagnose issues within the database. Several of the new statistics introduced in 10g were added specifically to improve the accuracy of the automated diagnosis of performance issues. One advantage of producing a tool inside the server is that if a problem is hard to diagnose we can add more instrumentation to make it easier!

6.1.1 Wait Classes

There are now over 700 different wait events possible in an Oracle Database 10g. The main reason for the increase is that many of the locks and latches have been broken out as separate wait events to allow for more accurate problem diagnosis. To enable easier high-level analysis of the wait events they have been categorized into WAIT CLASSES based on the solution space that normally applies to fixing a problem with the wait event. For example exclusive TX locks are generally an application level issue and HW locks are generally a configuration issue. The most commonly occurring wait classes and a few examples are listed below:

1. Application - locks waits caused by row level locking or explicit lock commands
2. Administration – DBA commands that cause other users to wait like index rebuild
3. Commit – wait for redo log write confirmation after a commit
4. Concurrency – concurrent parsing and buffer cache latch and lock contention
5. Configuration –undersized log buffer space, log file sizes, buffer cache size, shared pool size, ITL allocation, HW enqueue contention, ST enqueue contention
6. User I/O – wait for blocks to be read off disk
7. Network Communications – waits for data to be sent over the network
8. Idle – wait events that signify the session is inactive such as 'SQL*Net message from client'

6.2 AWR: A Repository of Performance Information

The importance of maintaining a repository of information about the operation of an Oracle system is achieved in Oracle 10g by AWR. The earlier Statspack scripts shipped with the database since 8i have been very widely used. In 10g this has been taken to the next level with the introduction of the Automatic Workload Repository. AWR runs automatically to collect data about the operation of the Oracle system and stores the data that it captures into the database. AWR is designed to be lightweight and to self manage its use of storage space so that you don't end up with a repository of performance data that is larger than the database that it is capturing data about.! After a default installation AWR will capture data every 30 minutes and will purge data that is over 7 days old. Both the frequency and length of time for which data is kept can be configured. Manual snapshots can also be performed.

Most database will by default be configured to collect data every hour upon the hour, and this has been found to be more than sufficient for most installations.

AWR captures all of the data previously captured by Statspack plus the new data described above. The data captured allows both system level and user level analysis to be performed, again reducing the requirement to repeat the workload in order to diagnose problems. Optimizations have been performed to ensure that the capture of data is performed efficiently to minimize

overhead. One example of these optimizations is in the SQL statement capture. Working within the database are maintained deltas of the data for SQL statements between snapshots. These allow the capture only of statements that have significantly impacted the load of the system (across a number of different dimensions such as CPU and elapsed time,) since the previous snapshot in an efficient manner, rather than having to capture all statements that had performed above a threshold level of work since they first appeared in the system, as was previously the case. This both improves the performance of the SQL capture and greatly reduces the number of SQL statements that are captured over time. Statements are captured based on the cumulative impact of all executions over the time period, so a heavily executed statement that completes in less than one second per execute will be captured alongside a single parallel query that ran for 15 minutes. The new Active Session History (ASH) data is captured in to AWR by 'sampling the samples' to reduce the volume of data. ASH data may also be flushed into the AWR on disk structures between snapshots if its circular buffer fills up. The new statistics and workload repository provide the basis for improved performance diagnostic facilities in Oracle Database 10g which support both proactive and reactive monitoring with ADDM and the EM Performance Page respectively.

The AWR report output closely resembles that generated from STATPACK. For this reason STATSPACK reports have not been incorporated as part of the STATS package.

6.3 Automatic Database Diagnostic Monitor: Proactive Diagnostics

Building upon the data captured in AWR, the Oracle Database 10g includes the Automatic Database Diagnostic Monitor (ADDM), a holistic self-diagnostic engine built right into the database. Using a medical analogy, using ADDM is similar to visiting a General Practitioner. It looks at the whole system, gives a diagnosis and then either suggests treatment itself or it may refer to specialists, other 10g advisory components, such as the SQL tuning advisor. As ADDM runs automatically after each AWR statistics capture it could be thought of as regularly scheduled performance checkups. In much the same way that a doctor will treat one regardless of race or creed ADDM will be equally at home working on any type of database, OLTP, data warehouse or mixed.

The goal of ADDM is to identify those areas of the system that are consuming the most 'DB time'. ADDM drills down to identify the root cause of problems rather than just the symptoms and reports the impact that the problem is having on the system overall. If a recommendation is made it reports the benefits that can be expected, again in terms of time. The use of time throughout allows the impact of several problems or recommendations to be compared. Previously many problems have been identified based on value judgments and experience rather than quantifiable impacts. A good example of this is a system that is experiencing a high logon rate. A rule of thumb might have said that a logon rate of greater than 10 per seconds was a problem and should be fixed. However many systems can run significantly higher logon rates without it noticeably affecting performance.

ADDM also documents the non-problem areas of the system. Wait classes that are not significantly impacting the performance of the system are pruned from the classification tree at an early stage and are listed so that the DBA can quickly see that there is little to be gained by performing actions in those areas. Again this saves time and wasted effort (both human and hardware) fixing things that will not impact the system performance overall. The report produced by ADDM are available in both a textual report and though EM. Using EM the ADDM findings are available right from the Database Home Page. Clicking on the link will take one to the ADDM Findings Screen and then to any recommendations.

Performance Collection

The STATS package has been extensively upgraded and documented since its original adhoc creation. New functionality has been added to enable its usage upon any version of Oracle database. The code is documented clearly with reference to changes that are required to enable successful compilation upon all databases. Where possible dynamic SQL has been employed to avoid the need to make changes for each individual installation.

7 Site Specific Parameters

There are some very specific site changes that need to be specified when the package is installed. A few of these are listed below. See Appendix A for full details.

- Output file directory: The location within the file system where output reports are to be generated. The directory name 'LOGDIR' is created as part of the installation and the user is prompted for the specification to be used. This is also used in calling scripts if the procedure in the package is not used.
- Mailserver: The fully qualified domain name (or the IP address) of the SMTP mail server used for automatic email generation.
- Email recipient The name of the person who will receive email.
- Email sender The name that the Oracle database will use to indicate as the originator of the message. This identifier must be one that the SMTP server will recognise and handle mail for.

See Appendix A for more specific instructions on parameter values and the installation process.

7.1 Email problems

There may be circumstances where apparently despite the correct settings being provided for the SMTP server that email is still being rejected. The most common of these problems relate to the SMTP server refusing to RELAY the message. This is always a configuration issue with the SMTP server. It is not unusual for SMTP servers (particularly Exchange) to have relaying turned off. It is possible to configure the server to accept messages coming from a supplied list of IP addresses. See your friendly SMTP (or Exchange) system administrator.

8 Package Descriptions

This section describes the calling parameters for every procedure and function used within the package. The external procedures are the ones that will be used in most usual day to day activities. The internal procedures are provided for completeness but are not callable except from within the package itself.

8.1 STATS Package

8.1.1 External procedures and functions

Procedure Name	Parameters	Param Type	Default value	Description
Debugging procedures				
PROCEDURE package_version	major_version	OUT NUMBER		Returns version information of the package.
	minor_version	OUT NUMBER		
PROCEDURE set_dbg	NONE	N/A	-	Used to allow the generation of informational messages intended to assist debugging during development.
PROCEDURE set_nodbg	NONE	N/A	-	As above but turns off the message generation.
FUNCTION debugging	NONE	N/A		RETURN BOOLEAN Generally only of use internally, determines the state of the internal debug flag.
Generic Utility Routines				
PROCEDURE execDDL	ddl_string	IN VARCHAR2		
PROCEDURE gen_schema				
PROCEDURE list_one	stat_type	IN VARCHAR2		
Statistics determination routines				
PROCEDURE db_space_critical_segs	no_extents	IN INTEGER	5	List any segments whose next n (5) extents cannot fit in the sum of all the free space left in a tablespace.
PROCEDURE db_space_critical_segsa	no_extents	IN INTEGER	5	Variation of the above only for archived tablespaces.
PROCEDURE io_spread				Displays IO figures for the database files and tablespaces.
PROCEDURE db_logs	chk_days	IN INTEGER	2	Displays the time between the last 'chk_days' log switches.
PROCEDURE list_sess_io				Lists session IO figures for all connected users when the report is run.
PROCEDURE alist_sess_io				Alternative display of the above.

Procedure Name	Parameters	Param Type	Default value	Description
PROCEDURE db_pins				Shows objects specifically pinned into the SGA.
PROCEDURE new_objects				Displays any database objects recently created.
PROCEDURE db_license				Displays the current database license settings
PROCEDURE file_limits				Indicates the limitations of the current database file settings
PROCEDURE db_io_disk				I/O distributions across disk files.
PROCEDURE chk_backups				Shows when the last database backups have been performed.
PROCEDURE temp_written				Amount of temporary data written.
PROCEDURE temp_usage				Current temporary data usage.
PROCEDURE temp_specs				Temporary data tablespace definitions
PROCEDURE wait_stats				Wait statistics.
PROCEDURE sort_params				Sort parameter settings
PROCEDURE check_db				Start-up and database general settings
PROCEDURE db_start				Displays startup time of database
PROCEDURE chained_rows				No of chained database objects.
PROCEDURE bcache				Buffer Cache ratio
PROCEDURE redo_trans				Redo statistics
PROCEDURE redo_groups				Redo group statistics
PROCEDURE roll_segs_accum				Rollback segments accumulated
PROCEDURE rbs_stats				Specific rollback segment statistics
PROCEDURE rbs_seg_stats				A variation upon a theme.
PROCEDURE extents	no_extents_in	IN INTEGER	10	Displays database objects with more than the specific number of extents
PROCEDURE ts_free_space				The amount of free space in each tablespace
PROCEDURE chk_enqueue_stats				Enqueue statistics
PROCEDURE io_reqs				Database I/O requests
PROCEDURE tab_growth	day_to_run	IN INTEGER	6	Updates tables with current tablespace usage.
PROCEDURE stat_run				Main calling routine that invokes a predetermined set of the above routines.
PROCEDURE stat_run_nosave				As 'stat_run' above only the run statistics are not saved.
PROCEDURE set_log_file	fname	IN VARCHAR2		Specifies the name of the output log file.
FUNCTION run_addm	start_time	IN DATE		RETURN VARCHAR2
	end_time	IN DATE		

Procedure Name	Parameters	Param Type	Default value	Description
PROCEDURE daily_run	P_options	IN INTEGER	0	Main daily run procedure P_options value are as follows: 0 Stats and ADDM outputs generated. 1 Stats, ADDM and AWR output in text format. 2 Stats, ADDM and AWR output in html format.
PROCEDURE create_daily_scheduled_job	P_options	IN INTEGER	0	Creates scheduled job for daily_run. This includes program, schedule and job. P_options value are as follows: 0. Stats and ADDM outputs generated. 1. Stats, ADDM and AWR output in text format. 2. Stats, ADDM and AWR output in html format.
PROCEDURE remove_daily_scheduled_job				Removed schedule job for daily run. This includes program, schedule and job.

Table 1 – Stats External procedures/functions

All of the above routines may be invoked separately, if there is a requirement for specific areas of investigation.

8.2 LOG Package

8.2.1 External procedures and functions

Procedure Name	Parameters	Param Type	Default value	Description
Debugging procedures				
PROCEDURE package_version	major_version	OUT NUMBER		Returns version information of the package.
	minor_version	OUT NUMBER		
PROCEDURE set_dbg	NONE	N/A		Used to allow the generation of informational messages intended to assist debugging during development.
PROCEDURE set_nodbg	NONE	N/A		As above but turns off the message generation.
FUNCTION debugging	NONE	N/A		RETURN BOOLEAN Generally only of use internally, determines the state of the internal debug flag.
PROCEDURE set_verbose		N/A		
PROCEDURE set_noverbose		N/A		
FUNCTION verbose		N/A		RETURN BOOLEAN
PROCEDURE set_log		N/A		

Procedure Name	Parameters	Param Type	Default value	Description
PROCEDURE set_nolog		N/A		
FUNCTION logging		N/A		RETURN BOOLEAN
Log File control Routines				
PROCEDURE Set_File_Separator	sep	IN VARCHAR2		Permits the specification of a file directory character if this can not be determined internally
PROCEDURE Create_Log	filename	IN VARCHAR2	LogFile	Create the specified log file. Default is the value of the LogFile specification in the Package Specification Header.
PROCEDURE Close_Log		N/A		Closes the current Log File
PROCEDURE log_mess	message_in	IN VARCHAR2		Writes the specified message to the specified log file.
	log_file	IN VARCHAR2	LogFile	
	Dbg	IN BOOLEAN	TRUE	Controls where the output is written to the screen and the logfile (TRUE set) or to the logfile only (FALSE)
PROCEDURE tslog_mess	message_in	IN VARCHAR2		Writes the specified message to the specified log file with an initial timestamp upon each output line.
	log_file	IN VARCHAR2	LogFile	
	Dbg	IN BOOLEAN	TRUE	Controls where the output is written to the screen and the logfile (TRUE set) or to the logfile only (FALSE)

Table 2 – Logging External procedures/functions

8.3 EMAIL Package

8.3.1 External procedures and functions

Procedure Name	Parameters	Param Type	Default value	Description
Debugging procedures				
PROCEDURE package_version	major_version	OUT NUMBER		Returns version information of the package.
	minor_version	OUT NUMBER		
PROCEDURE set_dbg	NONE	N/A	-	Used to allow the generation of informational messages intended to assist debugging during development.
PROCEDURE set_nodbg	NONE	N/A	-	As above but turns off the message generation.
FUNCTION debugging	NONE	N/A		RETURN BOOLEAN Generally only of use internally, determines the state of the internal debug flag.

Procedure Name	Parameters	Param Type	Default value	Description
Mailing Package Specific Routines				
PROCEDURE db_send_email	v_body	IN VARCHAR2	-	Test email generation routine.
	v_body2	IN VARCHAR2	-	
	v_email_address	IN VARCHAR2	-	
	v_subject	IN VARCHAR2	-	
	v_sender_address	IN VARCHAR2	-	
PROCEDURE mail_files	from_name	IN VARCHAR2		Required specifications are from_name, to_name, subject and message. Optionally up to three filenames may be specified for inclusion with the generated email message.
	to_name	IN VARCHAR2		
	subject	IN VARCHAR2		
	message	IN VARCHAR2		
	max_size	IN NUMBER	9999999999	
	filename1	IN VARCHAR2	NULL	
	Filename2	IN VARCHAR2	NULL	
	Filename3	IN VARCHAR2	NULL	

Table 3 – Email External procedures/functions

8.4 STATS_GEN Package

This package is created during the installation and is used within the installation script itself. It is not used at any other time and could be removed once installation is complete. It is recommended that it is retained, just in case there is ever a need to re-initialise the base tables.

See the installation script, or the specification header for the database package for details if required.

Using the Reports

9 Common performance areas

Below are typical events which frequently appear in the reports along with the relevant area to examine:

'db file scattered read' and 'db file sequential read' (and other IO related events)

The 'db file scattered read' and 'db file sequential read' are the two most common Read events Oracle waits for; db file scattered read indicates a full table scan is occurring, or waits for the db file sequential read event which indicates a single block read is occurring (Which one it waits for depends on the optimizer's determination of the best way to return the requested data).

The appearance of these events may not necessarily indicate a problem, as IO is a normal activity on a healthy instance. However, they can indicate problems if any of the following circumstances are true:

- The data-access method is bad (that is, the SQL statements are poorly tuned), resulting in unnecessary or inefficient IO operations
- The IO system is overloaded and performing poorly
- The IO system is under-configured for the load
- IO operations are taking too long

The above are usually tightly integrated.

To determine whether IO is an issue, examine the OS IO statistics (as described earlier) for symptoms, and compare with average time per read in the File and Tablespace IO sections of the reports. If the average time per read in the IO sections is large, and OS statistics indicate high service times or queue lengths, there is an IO problem.

Examine the SQL ordered by physical reads section of the report to see if there are any candidate high-resource SQL statements which can be tuned to reduce the IO load.

Tune these statements; tuning high-resource or frequently executed SQL can greatly reduce the IO load.

If the IO system continues to be overloaded, or the read times are still high, examine the host hardware for disk bottlenecks and identify how the files and/or disks can be reconfigured to spread the IO load. Further evidence of an IO bandwidth problem is the appearance of other IO related wait events (e.g. 'db file parallel write', 'direct read', 'direct write', and 'log file parallel write').

'latch free'

A latch is a low level resource used for protecting internal Oracle memory structures. A wait for a 'latch free' occurs when a server requests a latch and is unable to immediately acquire that latch. If latch free waits are high on the list, look at the Latch-specific sections to see which latches are contended for.

'enqueue'

An enqueue is another term for a lock. Locks protect shared resources and allow access to that resource via a queuing mechanism. Lots of time spent waiting for the 'enqueue' event can be caused by various problems. Look at the Enqueue Statistics section to identify which are the highest contended enqueues.

'free buffer waits'

A free buffer wait event occurs when a server would like a buffer, but there are no unused buffers immediately available. If the time spent for 'free buffer waits' is significant, this can either imply the buffer cache is too small, or that DBWR is not writing enough buffers to disk fast enough to keep up with requests for new buffers. Use O/S monitor tools to check the IO system, and look at the File IO statistics to examine whether the IO system may be slowing down DBWR.

'buffer busy wait'

A buffer busy wait event occurs when a server process would like to access a buffer which is either currently being read into the cache, or is already in the cache but is already being used in an unsharable way. Check the Buffer Wait Statistics section to identify the contended-for buffer types, and correlate this data with the wait data in the Tablespace and File IO sections to see whether there are any specific files or tablespaces which are experiencing buffer contention more than any others. This is the first step towards find out which segments are contended for, and why.

'write complete waits'

A write complete wait event occurs when DBWR is writing a block out to disk when a server process would like to use the block; this implies DBWR is too slow (and hence there is an IO bottleneck), or the cache is too small, or there is a number of processes performing a large numbers of indexed buffer gets. Take note of this event if it occurs between checkpoints (this event is normal during a checkpoint and can be ignored). To identify whether a SQL statement is causing large numbers of indexed buffer gets, examine the SQL section ordered by Buffer Gets to identify statements which may not be using the most selective indexes; using non-selective indexes unnecessarily flushes useful buffers from the buffer cache.

9.1 Important Database Ratio's

The following list explains how some of the important database ratio's are calculated and refers you to other areas of the report for investigating suspicious values. Although the calculations are actually percentages, the term ratio is used to be consistent with the report headings.

Buffer Nowait Ratio

Is the percentage of requests a server process makes for a specific buffer where the buffer was immediately available; all buffer types are included in this statistic. If the ratio is low, determine which type of block is being contended for by examining the Buffer Wait Statistics section of the report.

Buffer Hit Ratio

This statistic is also known as the buffer cache hit ratio. This is the percentage of requests for a particular block which are satisfied within the cache without the need for physical IO.

Although historically known as one of the most important statistics, the buffer cache hit ratio can sometimes be misleading. A high (e.g. 99%) cache hit ratio normally indicates the cache is adequately sized, however this may not be the case in all circumstances. For example, frequently executed SQL statements which repeatedly refer to a small number of buffers via indexed lookups can skew the buffer gets statistic. When these blocks are read, they are placed at the most recently used (MRU) end of the buffer cache; iterative access to these blocks can artificially inflate the cache hit ratio. This makes tuning the buffer cache a challenging activity.

Sometimes it is possible to identify a too small buffer cache by the appearance of the 'write complete waits' event, which indicates that hot blocks (i.e. blocks which are still being modified)

are aging out of the cache while they are still needed; check the Wait events section for evidence of this event.

Alternatively, a lower buffer cache hit ratio does not necessarily mean the cache is too small; it may be that (potentially valid) full table scans are artificially reducing what is otherwise a good hit ratio.

Library Hit Ratio

This is also known as the library cache hit ratio. The ratio indicates the number of pin requests which result in pin hits. A pin hit occurs when the SQL or PL/SQL code you wish to execute is already in the library cache and is valid to execute.

A low library cache hit percentage could imply SQL is prematurely aging out of the shared pool as the shared pool may be small, or that unsharable SQL is being used. Also compare with the soft parse ratio; if they are both low, then investigate whether there is a parsing issue.

Redo no-wait Ratio

This ratio is indicative of the number of redo-entries generated for which there was space immediately available in the redo log. The percentage is calculated as followed:

```
100 x (1- (redo log space requests/redo entries))
```

The 'redo log space request' statistic is incremented when an Oracle process attempts to write a redo entry, however there was not sufficient space remaining in the online redo log. The 'redo entries' statistic is incremented for each entry made to the redo log.

Frequent, or slow log switches may be contributing to waits for redo log space. If logs are switching frequently (e.g. more than once every 15 minutes) this may be improved by increasing the size of the online redo logs.

If the log switches are not frequent, check the disks the redo logs reside on to see if log switches are taking a long time due to a slow IO system. If the IO system is overloaded, either move the redo logs to disks with less activity, place the logs on dedicated disks or faster devices.

In-memory Sort Ratio

This is the percentage of sorts which were performed in memory, rather than sorts which also required a disk sort segment to complete. Optimally, in an OLTP environment the percentage of sorts performed in memory should be high; refer to the Oracle Manual Designing and Tuning for Performance manual (i.e. the "server tuning" guide) for information on tuning sorts.

Soft parse ratio

The soft parse ratio shows the total number of parses which were soft.

A soft parse occurs when a session attempts to execute a SQL statement, the statement is already in the shared pool, and can be used. For a statement to be used (i.e. shared) all data, (including data such as the optimizer execution plan) pertaining to the existing SQL statement must be equally applicable to the current statement being issued.

A hard parse occurs when a SQL statement is executed, and the SQL statement is either not in the shared pool, or it is in the shared pool but can not be shared as part of the metadata for the two SQL statements is different (for example, this may happen if a SQL statement is textually identical as a pre-existing SQL statement, but the tables referred to in the two statements resolve to physically different tables).

In an OLTP environment, hard parses are expensive CPU wise, which adds elapsed time to the user executing the statement. The aim is to parse once, execute many times. Ideally the soft parse ratio would be greater than 95%; when the soft parse ratio falls significantly lower than 80%, it may be cause to investigate whether it is possible to share SQL by using bind variables, or if the code can not be changed, to force cursor sharing by using the init.ora parameter introduced in Oracle8i release 8.1.6, cursor_sharing.

As a sanity check, compare this ratio to the hard and soft parse rates (per second) in the Load Profile. If the rates are low (e.g. 1 per second), parsing may not be a significant issue.

Another useful comparison is against the proportion of parse time that was not CPU-related:

```
(parse time CPU) / (parse time elapsed)
```

A low value for this ratio could mean that the non-CPU-related parse time was spent waiting for latches, which might indicate a parsing or latching problem. To investigate further, look at the shared-pool and library-cache latches in the Latch sections for indications of contention on these latches.

Latch Hit Ratio

This percentage is based on the ratio of the total number of latch misses to the number of latch gets for all latches. The ratio is indicative of a latching problem if the ratio is low, however as the data is rolled up over all latches, a high ratio can artificially mask a low get rate on a specific latch. Cross check this value with the top-5 wait events to see if 'latch free' is in the list, and if so, refer to the Latch sections.

.

A. Installation steps

As the 'SYS' user run the script 'load_code.sql. (See Appendix D for listing) This will prompt the user for the name of the owner of the code.

NOTE: If the user already exists they will be dropped.

The user is also prompted for the 'code owners' default tablespace.

A further prompt requests the full path where the output logs will be generated. A fully qualified name is required. i.e. C:\TEMP or something similar if on Microsoft Windows. This will create a database object called 'LOGDIR'.

Finally the script will prompt for the code owner's passwd.

The script creates all required database objects, and provides all relevant roles and privileges.

A.1 Local changes

There are two ways in which local changes can be specified. The preferred way is to update the values in the table PACKAGE_PARMS.

Table 4 - PACKAGE_PARMS Table

PNAME	PARAMNAME	VALUE (Default)	Description
EMAIL	MAILHOST	mail.somewhere.com	
EMAIL	MAILPORT	25	
LOGGING	LOGFILE	C:\TEMP\file.log	
STATS	OUTPUT_NAME	C:\TEMP\new.lst	
STATS	RECIPIENT	anon@somewhere.com	
STATS	SENDER	anon@somewhere.com	
STATS	REPORT_TITLE	Stats and ADDM Report	
STATS	AWR_LOG_FILE	C:\TEMP\awr_report.html	
STATS	AWR_TLOG_FILE	C:\TEMP\awr_report.txt	

The table PACKAGE_PARMS is created as part of the installation and is populated with the default values. The site must modify the defaults to suitable values appropriate for the site.

A simple update statement of the following form can be used:

```
UPDATE PACKAGE_PARMS
Set VALUE='xxxxx'
WHERE PNAME='yyyyy'
AND PARAMNAME='zzzzz';
COMMIT;
```

Where xxxxx is the new value, and yyyyy is the Package name, and zzzzz is the paramname from the above table. Failure to set the values appropriately will result in the programs failing to run successfully.

AN alternative can be used where the PACKAGE_PARMS table does not exist in which case the values hard coded in the package specification headers will be used. If a specific parameter is missing from

the PACKAGE_PARMS table the specification header value will be used. A description of the parameters and their usage follows.

A.1.1 ADDM Settings

The STATS package parameters are as follows:

output_name VARCHAR2(30) DEFAULT 'C:\TEMP\new.lst';

The desired output file name for the daily report runs. This file will be removed by the process and recreated every time the schedule is run

recipient VARCHAR2(30) DEFAULT 'anon@somewhere.com';

The name of the recipient of the generated report emailed to the user. Only one person is currently set up to receive email. If more than one is desired, set up an address list on the SMTP server and set this value to the name of the address list.

sender VARCHAR2(30) DEFAULT 'anon@somewhere.com';

The name by which the Database is known to the SMTP server. It is usual for the SMTP server to check the sender names before deciding whether to act as a relay for the message. If there are problems with the email being received by the recipient check the SMTP server logs to see if the message has been rejected, and whether relaying is enabled.

report_title VARCHAR2(30) DEFAULT 'Stats and ADDM Report';

A.1.2 AWR Settings

If the AWR reports are to be run a few further steps need to be performed.

Change the default values for the STATS package as follows:

awr_log_file VARCHAR2(30) := 'C:\TEMP\awr_report.html';

Use to generate the HTML formatted AWR report.

awr_tlog_file VARCHAR2(30) := 'C:\TEMP\awr_report.txt';

This parameter specifies the log file name for the text version of the AWR report.

A decision need to be made as to whether the text or the html versions of the report are to be used. The default is to generate the html version.

A.1.3 DBMS schedule setting

Create the daily schedule by running the following:

```
Begin
Stats.create_daily_scheduled_job(Option);
End;
/
```

Where *Option* has a value of 0, 1 or 2. Default if no value is specified is 0.

The values are as follows:

0 Generate Stats and ADDM report.

1 Generate Stats, ADDM report and also AWR report in 'text' format for the same period.

2 Generate Stats, ADDM report and also AWR report in 'html' format for the same period.

A.1.4 Email settings

The EMAIL package parameters are as follows:

MAILHOST VARCHAR2(30) := 'mail.somewhere.com';

The name of the SMTP server with will act as a relay for the email message. For most installations this would be the address of the Exchange server if Microsoft Exchange is being used on the site.

MAILPORT INTEGER := 25;

The port number upon which the SMTP server listens for incoming mail requests. This value should not need to be changed except under very unusual site configurations.

A.1.5 Logging Settings

This parameter is used where no output file specification is specified elsewhere. It is really a catch all and will not be used if the installation steps are followed correctly.

A.2 Installation Testing

The following steps test out the installation and assist in tracking down any problems that may be encountered.

A.2.1 Test of basic stats installation

Run the script in B.1 to check that an output is generated in the specified directory.

An output file should be generated. If the output file is not generated, check for any reported errors. Using the 'set serverout on' specification in SQL*Plus may provide some additional information. The most likely cause of problems is the specifications of the 'LOGDIR' directory and the specification of the log file not matching. The 'case' of the specification can often be a problem.

As the 'STATS' package owner run the following SQL

SQL> select * from all_directories;

Look for the 'LOGDIR' specification and modify if incorrect, or specify the correct values in the PACKAGE_PARMS table.

This test also checks that the LOGGING package is working correctly.

A.2.2 Further check on STATS outputs

If the test above is successful run the script in B.2 to check that all outputs are generated. Specify options (1 and 2) to confirm that both AWR reports are generated. The ADDM report will be appended to the bottom of the main stats log file.

A.2.3 Check of email installation.

Run the script given in B.3, using suitable values for the sender, recipient, and filename1. This test just checks that email can be send from the system. It is a separate component and works in its own right.

If email is not generated, check the settings for the EMAIL MAILHOST specified in the PACKAGE_PARMS table. [An alternative would be to check the value in the EMAIL package specification. Note that a value in the PACKAGE_PARMS table will override and package header specification value.]

Running with 'set serverout on' may provide some additional information.

If still unsuccessful, check the SMTP server to ensure that it is accepting (and logging) the input. Any further problem areas are most likely to be at the SMTP server end, rather than with the EMAIL package itself, especially if the package does not generate any error information.

A.2.4 Putting it all together

Edit the script provided in B.5 with suitable values and run the reports. As shown the script will generate a report covering the past two hours. Note that if a system restart has occurred with the specified time period an error will be generated. In this situation run two reports around the restart time.

If all the previous tests have run successfully then the statistics report will be emailed to the recipient.

A.2.5 Testing the scheduler

It is slightly more difficult to test the schedule, since the schedule is set up to run at 8 pm every evening, and report over a time period of from 7am to 7pm.

Having run the scheduler set up job (See A.1.3) inspection of the various scheduler tables can be used to check the values.

Looking at the user_scheduler_programs, user_scheduler_schedules and the user_scheduler_jobs tables can assist in tracking down problems.

Whilst the provided set up has been tested and proved to work as specified, occasional problems may be encountered if the database is not set up correctly. Please see the system DBA for further assistance.

B. Sample Scripts

B.1 Script to emulate traditional script report

```
BEGIN
  stats.set_log_file('C:\TEMP\stats.log');
  stats.stat_run;
END;
```

B.2 Script to generate typical run.

This example is as used internally by the Stats package to generate the daily runs. An optional parameter can be specified, which can be a value of 0, 1 or 2. See section 8.0 for details of the meaning of the parameters.

```
BEGIN
  Stats.daily_run;
END;
/
```

B.3 Script to test email installation

```
BEGIN
  email.mail_files(
     to_name => 'anon@somesite.com',
     from_name => 'dba@somesite.com',
     subject => 'ADDM Report',
     message => 'Todays ADDM report',
     filename1 => 'C:\TEMP\stats.log');
END;
/
```

B.4 Full Script to display latest ADDM report

```
set long 1000000
set pagesize 50000
column get_clob format a80
select dbms_advisor.get_task_report(task_name) as ADDM_report
from dba_advisor_tasks
where task_id = (
select max(t.task_id)
from dba_advisor_tasks t, dba_advisor_log l
where t.task_id = l.task_id
and t.advisor_name = 'ADDM'
and l.status = 'COMPLETED');
```

B.5 Full Script to generate and email report

This is an example of a script that would send email. Modify to suit the particular site requirements.

```
-- set SQL*Plus variables and column formats for the report
spool today.lst
SET PAGESIZE 0 LONG 1000000 LONGCHUNKSIZE 1000;
COLUMN get_clob FORMAT a80;

VARIABLE task_name VARCHAR2(30);
declare
  yesterday  DATE;
```

```
  imonth     VARCHAR2(2);
  iday       VARCHAR2(2);
BEGIN

/*   The following code merely sets the start and end time of the required report
        Period.
*/
    -- yesterday := trunc(sysdate-1);
    yesterday := sysdate;

    imonth := to_char(yesterday,'MM');
    iday   := to_char(yesterday,'DD');

    dbms_output.put_line('Date: '||to_char(yesterday));
    /*
    :task_name := stats.run_addm( TO_DATE('07:00:00 '||iday||imonth||'', 'HH24:MI:SS
DDMM'),
                    TO_DATE('19:00:00 '||iday||imonth||'', 'HH24:MI:SS DDMM') );
    */
--  Generate a report covering the past two hours.
--  Any time periods can be specified as desired.
--
    :task_name := stats.run_addm( sysdate-2/24, sysdate);
END;
/
Rem  specifiy a desired output file name.
spool C:\temp\today.lst
-- execute GET_TASK_REPORT to get the textual ADDM report.
SELECT DBMS_ADVISOR.GET_TASK_REPORT(:task_name)
FROM DBA_ADVISOR_TASKS t
WHERE t.task_name = :task_name
AND t.owner = SYS_CONTEXT( 'userenv', 'session_user' );
spool off
BEGIN
-- Fill in the parameters with the required values.
  email.mail_files(
     to_name => 'anon@somesite.com',
     from_name => 'dba@somesite.com',
     subject => 'ADDM Report',
     message => 'Todays ADDM report',
     filename1 => 'C:\TEMP\today.lst');
END;
/
```

C. Sample Outputs

C.1 Sample Output from traditional script run

[Note: This was run against a development database so not too much value should be placed upon the reported figures. It is shown for demonstration purposes only.]

```
    Database Performance/Information Report
    =======================================

 This script queries various Oracle Data Dictionary views to collect
 basic statistics data reflecting a database's configuration and prompt performance.

 Oracle Server is a sophisticated dynamic mechanism so no intention
 is to be made out of THIS SCRIPT other than to reveal a full
 (or un-misleading) picture of a targeted database's performance.

 An accurate assessment on the database relies on a lot more work involving:
 1. more sophisticated queries and OS commands in addition to this script
 2. continuous and proper monitoring and sampling; and
 3. most importantly a correct and coherent interpretation of any statistics
    data collected in conjunction with the database configuration and
    applications in question.

Db Check Starts: (DD-MON-RR HH24:MI:SS)
---------------------------------------
01-FEB-05 16:49:01


 ******* The database to be checked:

Name      Created         Log Mode     Checkpoint Change Archive Change
--------- --------------- ------------ ----------------- -----------------
GEN2      04-MAY-04 14:15 NOARCHIVELOG     3031579402879     3031579345396

Instance startup time
---------------------
01-FEB-2005 08:02


 1. Check the Shared Pool
 ========================

 1.1 The shared SQL area
 =======================


 1.1.1  Are cursors (or SQL statements) being shared?
 Hit Ratio in Library Cache should be > 95%. If not, there is probably room
 to improve the efficiency of the APPLICATION CODE.

Hit ratio %
-----------
      98.54


Activity                                            Log time          Value
--------------------------------------------------- ----------------- ----------
shared sql cache hit ratio (should be > 95%)        01-FEB-2005 14:45      92.93
                                                    01-FEB-2005 15:26      97.15
                                                    01-FEB-2005 16:41      98.48
                                                    01-FEB-2005 16:49      98.54

 1.1.2  Is there enough space to cache cursors (or SQL stmts)?

 PINS:    executions of an item in the library cache.
 RELOADS: an item loaded for this cursor has been  aged out for lack of space prompt
 The Reloads / Pins should be < 1%. Otherwise increase SHARED_POOL_SIZE.

Reload ratio %
--------------
           .09


Activity                                            Log time          Value
--------------------------------------------------- ----------------- ----------
shared sql cache reload ratio (should be < 1%)      01-FEB-2005 14:45        .15
                                                    01-FEB-2005 15:26        .11
                                                    01-FEB-2005 16:41        .09
                                                    01-FEB-2005 16:49        .09

 1.1.3  For all usage of the library cache.

        Reload ratio is sometimes also known as the Miss Ratio.

  Executions       Cache misses    Library Cache
                  while executing     Reload Ratio
 -----------    ---------------   --------------
     697,141               592                0


 1.1.4  Look at Library Cache in detail.

 Hit ratio should be >70
 Pin ratio should be >70 . . .
```

```
Namespace         Hit ration %  Pin hit ration %        Reloads
----------------- ------------  ---------------- --------------
SQL AREA                    98               99            562
TABLE/PROCEDURE             73               92              2
BODY                        99               99             28
TRIGGER                     95               99              0
INDEX                       17               41              0
CLUSTER                     96               98              0
OBJECT                     100              100              0
PIPE                       100              100              0
JAVA SOURCE                100              100              0
JAVA RESOURCE              100              100              0
JAVA DATA                   94               50              0
```

1.1.5 Look at Object Pinning.

```
Owner       Object Type  Object Name                      Executions  Loads  Memory Used Kept?
----------- ------------ -------------------------------- ---------- ------ ------------ ----
SYS         PACKAGE BODY DBMS_OUTPUT                          44,275      1        6,263 NO
SYS         PACKAGE BODY UTL_FILE                             41,443      1       19,216 NO
SYS         PACKAGE BODY STANDARD                             10,315      1       22,700 NO
GSC         PACKAGE BODY LOG                                   3,104     12       17,779 NO
SYS         PACKAGE BODY DBMS_APPLICATION_INF                  2,177      1        3,485 NO
SYSMAN      PACKAGE BODY MGMT_LOG                              2,169      1        6,819 NO
SYSMAN      TRIGGER      RAW_METRICS_AFTER_IN                  1,315      1        5,831 NO
SYSMAN      PACKAGE BODY EMD_SCHEMA                            1,315      1       17,229 NO
SYSMAN      PACKAGE BODY MGMT_FAILOVER                         1,028      1        9,004 NO
SYSMAN      PACKAGE BODY EMD_LOADER                              568      1       72,205 NO
SYSMAN      PACKAGE BODY EMD_MAINTENANCE                         562      1       71,570 NO
SYSMAN      PACKAGE BODY MGMT_JOB_ENGINE                         515      1      248,314 NO
SYSMAN      PACKAGE BODY EMD_COLLECTION                          514      1       34,969 NO
SYSMAN      PACKAGE BODY EM_SEVERITY_REPOS                       514      1       30,600 NO
SYSMAN      PACKAGE BODY MGMT_GLOBAL                             514      1        8,858 NO
SYSMAN      PACKAGE BODY EM_PING                                 514      1       32,806 NO
SYS         PACKAGE BODY DBMS_STATS_INTERNAL                     499      1       73,679 NO
SYS         PACKAGE BODY DBMS_RCVMAN                             490      1      364,011 NO
```

1.2 The data dictionary cache

```
GETS:      requests for information
GETMISSES: requests resulting in cache misses
Hit  Ratio should be > 90%
Miss Ratio should be < 15%
Otherwise consider increasing SHARED_POOL_SIZE.
```

```
Miss ratio %
------------
         .99
```

```
Activity                                            Log time          Value
--------------------------------------------------- ----------------- ----------
data dictionary miss ratio (should be < 15%)        01-FEB-2005 14:45       9.87
                                                    01-FEB-2005 15:26       1.30
                                                    01-FEB-2005 16:41       1.00
                                                    01-FEB-2005 16:49        .99
```

2. The Buffer Cache

```
db block gets:   accesses to current copies of blocks (from cache or disk)
consistent gets: # of blocks read from read-consistent blocks
physical reads:  # of blocks read from disk
```

```
Name                                     Value
---------------------------------- ------------
db block gets                          218441
consistent gets                       3726377
physical reads                         371175
```

2.1 Hit Ratio:

```
Hit Ratio should be > 90%. Otherwise adjust DB_BLOCK_BUFFERS.
```

```
Hit Ratio %
-----------
      90.59
```

```
Activity                                            Log time          Value
--------------------------------------------------- ----------------- ----------
buffer hit ratio (should be > 90%)                  01-FEB-2005 14:45      94.07
                                                    01-FEB-2005 15:26      88.84
                                                    01-FEB-2005 16:41      90.51
                                                    01-FEB-2005 16:49      90.59
```

2.2 Other Buffer performance indicators

```
If the values are high, or much increased against previous ones,
consider prompt increasing the buffer cache.

free buffer inspected: the number of buffers skipped before finding a free buffer
```

```
Free buffers inspected
----------------------
               373418
```

```
Activity                                              Log time          Value
----------------------------------------------------- ----------------- ----------
free buffers inspected (Any big increase?)            01-FEB-2005 14:45      20662
                                                      01-FEB-2005 15:26     138952
                                                      01-FEB-2005 16:41     356255
                                                      01-FEB-2005 16:49     373418

buffer busy waits: the number of waits for a buffer to become available

Event                Total Waits Total Timeouts Time waited Average Wait
-------------------- ----------- -------------- ----------- ------------
buffer busy waits              5              0          11            2


Activity                                              Log time          Value
----------------------------------------------------- ----------------- ----------
buffer busy waits (Any big increase?)                 01-FEB-2005 14:45          5
                                                      01-FEB-2005 15:26          5
                                                      01-FEB-2005 16:41          5
                                                      01-FEB-2005 16:49          5


Activity                                              Log time          Value
----------------------------------------------------- ----------------- ----------
buffer busy wait time (Any big increase?)             01-FEB-2005 14:45          2
                                                      01-FEB-2005 15:26          2
                                                      01-FEB-2005 16:41          2
                                                      01-FEB-2005 16:49          2


 3. The usage of Oracle Blocks
 ==============================
 3.1.  Are there any rows (in any tables) having chained rows?

No chained rows found in the database


 3.2  Any latch contention of "cache buffers chains"

 3.2.1. cache buffers chains: many waits on this latch typically indicate block contention.

 The latch hit ratio should be > 98%


   Latch#      Gets    Misses    Sleeps Hit Ratio %
--------- --------- --------- --------- -----------
      116   7236397         6         5      100.00


Activity                                              Log time          Value
----------------------------------------------------- ----------------- ----------
cache buffers chains hit ratio (Should be > 98%)      01-FEB-2005 14:45     100.00
                                                      01-FEB-2005 15:26     100.00
                                                      01-FEB-2005 16:41     100.00
                                                      01-FEB-2005 16:49     100.00

 To verify any contention above, look at the # of waits while accessing data blocks.

 Miss Ratio = Number of waits * 100 / Total number of data block gets

 The ratio should be <= 0.5%

Data Block Waits Miss Ratio %
---------------- ------------
             124          .00


Activity                                              Log time          Value
----------------------------------------------------- ----------------- ----------
access to data block miss ratio (Should be < 0.5%)    01-FEB-2005 14:45        .02
                                                      01-FEB-2005 15:26        .00
                                                      01-FEB-2005 16:41        .00
                                                      01-FEB-2005 16:49        .00


 3.2.2. cache buffers lru chain: high contention means a cache buffer
 is either too small or (less likely!) too big.

 The latch hit ratio should be > 98%

Name                                     Gets    Misses    Sleeps Hit Ratio %
------------------------------------ --------- --------- --------- -----------
cache buffers lru chain                 521843         0         0     100.000


Activity                                              Log time          Value
----------------------------------------------------- ----------------- ----------
cache buffers lru chain hit ratio (should be > 98% ? 01-FEB-2005 14:45     100.00
                                                      01-FEB-2005 15:26     100.00
                                                      01-FEB-2005 16:37     100.00
                                                      01-FEB-2005 16:41     100.00
                                                      01-FEB-2005 16:49     100.00

 3.3.  Any contention for freelist?

 NOTE: The "time" column below will be 0 if "timed_statistics" param is FALSE.


Operation class    # of waits by this operation total wait time
------------------ ---------------------------- ---------------
free list                                     0               0
```

Activity	Log time	Value
free list waits (Any big increase?)	01-FEB-2005 14:45	0
	01-FEB-2005 15:26	0
	01-FEB-2005 16:41	0
	01-FEB-2005 16:49	0

Activity	Log time	Value
free list wait time (Any big increase?)	01-FEB-2005 14:45	0
	01-FEB-2005 15:26	0
	01-FEB-2005 16:41	0
	01-FEB-2005 16:49	0

3.4. Any contention for segment header?

A high number might require more freelists, and check v$session_wait to get the addresses of the actual blocks having contention.

NOTE: The "time" column below will be 0 if "timed_statistics" param is FALSE.

Operation class	# of waits by this operation	total wait time
segment header	4	2

Activity	Log time	Value
segment header waits (Any big increase?)	01-FEB-2005 14:45	3
	01-FEB-2005 15:26	3
	01-FEB-2005 16:41	4
	01-FEB-2005 16:49	4

Activity	Log time	Value
segment header wait time (Any big increase?)	01-FEB-2005 14:45	1
	01-FEB-2005 15:26	1
	01-FEB-2005 16:37	2
	01-FEB-2005 16:41	2
	01-FEB-2005 16:49	2

3.5. Tablespace Free Space Information

Get the total freespace (in Mb and blocks) in each tablespaces:

Note the Number of Free Chunks in a tablespace.

Tablespace name	Free Bytes	Free	# of Free
PMS_DATA	983,04	120	1
PMS_IDX	655,36	80	2
SYSAUX	23,461,88	2,864	28
SYSTEM	16,252,92	1,984	5
UNDOTBS1	387,514,36	47,304	10
USERS	20,709,37	2,528	3
WIXBIL	21,561,34	2,632	2
sum		79,928	

3.6. Checking Fragmented Objects

Fragmentation Report: Fragmentation Report for Multiple Extents

Note the Number of Extents.
(only listing those segments with # of Extents > 10)

Object Name	Obj Type	Initial Extent	Next Extent	Pct Inc	Space Used Kb	# of Blks	# of Exts	Max Extents
MAG.CLONE_CAPTURE_P2	TABLE SUB	65536	0	0	2,108,160	263,52	32940	2147483645
MAG.CLONE_CAPTURE_P3A	TABLE PAR	65536	0	0	989,632	123,70	14758	2147483645
MAG.CLONE_CAPTURE_P1	TABLE PAR	65536	0	0	346,560	43,32	1815	2147483645
MAG.SYS_IOT_TOP_58715	INDEX PAR	65536	0	0	280,000	35,00	910	2147483645
MAG.SYS_IOT_TOP_58713	INDEX	65536	0	0	34,816	4,35	49	2147483645
WIXBIL.DAILY_MOVEMENT_IX2	INDEX	41943040	0	0	40,960	5,12	40	2147483645
WIXBIL.DAILY_MOVEMENT_IX3	INDEX	41943040	0	0	40,960	5,12	40	2147483645
WIXBIL.DAILY_MOVEMENT_IX4	INDEX	41943040	0	0	40,960	5,12	40	2147483645
WIXBIL.INVOICE_DETAIL_IX2	INDEX	36700160	0	0	35,840	4,48	35	2147483645
MAG.CLONE_CAPTURE	TABLE	65536	0	0	18,432	2,30	33	2147483645
CAI.SYS_LOB0000062415C00039$$	LOBSEGMEN	65536	0	0	10,240	1,28	25	2147483645
UPLOAD.VPCPRDEE	TABLE	65536	0	0	6,144	76	21	2147483645
UPLOAD.PRDMSTEE	TABLE	65536	0	0	5,120	64	20	2147483645
CAI.SYS_IMPORT_FULL_01	TABLE	65536	0	0	3,072	38	18	2147483645
GSC.WK$CACHE	TABLE	65536	0	0	3,072	38	18	2147483645
SYSMAN.MGMT_METRICS_RAW_PK	INDEX	65536	0	0	2,048	25	17	2147483645
CAI.EXPLAIN_STATS	TABLE	1048576	0	0	1,024	12	16	2147483645
CAI2.EXPLAIN_STATS	TABLE	1048576	0	0	1,024	12	16	2147483645
MAG.HOTLISTMATCH$IDX1	INDEX	2097152	131072	0	2,048	25	16	2147483645
MAG.HOTLISTMATCH$IDX2	INDEX	2097152	131072	0	2,048	25	16	2147483645
MAG.LASTCONTACT$IDX	INDEX	2097152	131072	0	2,048	25	16	2147483645
MAG.LASTCONTACT$IDX2	INDEX	2097152	131072	0	2,048	25	16	2147483645
MAG.PARAMETERS$IDX	INDEX	2097152	131072	0	2,048	25	16	2147483645
MAG.VPAMMATCH$IDX2	INDEX	2097152	131072	0	2,048	25	16	2147483645
MAG.VPAMMATCH$IDX3	INDEX	2097152	131072	0	2,048	25	16	2147483645
MAG.VPAMMATCH$IDX4	INDEX	2097152	131072	0	2,048	25	16	2147483645
MAG.VPAMMATCH$IDX5	INDEX	2097152	131072	0	2,048	25	16	2147483645
PMS.ACTIONS	TABLE	819200	0	0	896	11	14	2147483645
PMS.SYS_LOB0000061453C00002$$	LOBSEGMEN	131072	0	0	896	11	14	2147483645
GSC.DRWKDOC_PATH_IDX$I	TABLE	5242880	0	0	13,312	1,66	13	2147483645
WIXBIL.INVOICE_DETAIL_IX1	INDEX	83886080	0	0	81,920	10,24	11	2147483645

3.7 Check for Enqueue Waits
==========

```
enqueue waits Num transactions Per transaction Num logins Per logon Comments
------------- ---------------- --------------- ---------- --------- --------------------------
         8057             2858          2.8191        653   12.3384 increase enqueue_resources
```

4. Rollback Segments

4.1 General RBS statistics sizing

NOTE: check the sizing:
Check INITIAL_EXTENT and NEXT_EXTENT of the RBSs
The two value should be the same, and be 2 ^ N (power) bytes

RBS Name	Size in Bytes	Optimal Size	High Water Mark	# of Shrinks	# of Wraps	# of Extends	Mean size shrinks	Mean size active Extents
SYSTEM	385,024		385,024	0	0	0	0	0
_SYSSMU1$	3,268,608		6,414,336	1	10	2	5,242,880	315,449
_SYSSMU10$	253,952		253,952	0	0	0	0	0
_SYSSMU2$	4,317,184		9,560,064	1	11	3	8,388,608	388,761
_SYSSMU3$	2,220,032		2,220,032	0	7	0	0	272,200
_SYSSMU4$	3,268,608		3,268,608	0	2	0	0	199,228
_SYSSMU5$	4,186,112		4,186,112	0	0	0	0	0
_SYSSMU6$	2,220,032		2,220,032	0	0	0	0	0
_SYSSMU7$	2,351,104		2,351,104	0	0	0	0	0
_SYSSMU8$	385,024		385,024	0	0	0	0	0
_SYSSMU9$	253,952		253,952	0	0	0	0	0

4.2. Shows actual RollBack SEGMENT STATISTICS.....

Rollback Statistics

Name	Extents	RSSize	Number of Active tx	Size (Mb) written	Header gets	Header waits	Hit Ratio
SYSTEM	6	385024	0	.00	208	0	100.00
_SYSSMU1$	5	3268608	0	3.21	2331	1	99.96
_SYSSMU10$	4	253952	0	.00	221	0	100.00
_SYSSMU2$	6	4317184	0	4.36	3021	1	99.97
_SYSSMU3$	4	2220032	0	3.16	2365	0	100.00
_SYSSMU4$	5	3268608	0	1.82	5270	0	100.00
_SYSSMU5$	19	4186112	1	.76	631	0	100.00
_SYSSMU6$	4	2220032	0	.14	552	0	100.00
_SYSSMU7$	6	2351104	0	.11	340	0	100.00
_SYSSMU8$	6	385024	0	.00	221	0	100.00
_SYSSMU9$	4	253952	0	.00	221	0	100.00
sum				13.57			

*** The accumulated write size of Rollback Segs datafiles

Datafile name	Phy. Size written Mb

4.3 undo waits

NOTE: The "time" column below will be 0 if "timed_statistics" param is FALSE.

Operation class	# of waits by this operation	Total wait time
undo header	2	0

Activity	Log time	Value
undo header waits (Any big increase?)	01-FEB-2005 15:26	2
	01-FEB-2005 15:40	2
	01-FEB-2005 16:42	2
	01-FEB-2005 16:50	2

Activity	Log time	Value
undo header wait time (Any big increase?)	01-FEB-2005 15:26	0
	01-FEB-2005 15:40	0
	01-FEB-2005 16:42	0
	01-FEB-2005 16:50	0

Operation class	# of waits by this operation	Total wait time
undo block	0	0

Activity	Log time	Value
undo block waits (Any big increase?)	01-FEB-2005 15:26	0
	01-FEB-2005 15:40	0
	01-FEB-2005 16:42	0
	01-FEB-2005 16:50	0

Activity	Log time	Value
undo block wait time (Any big increase?)	01-FEB-2005 15:26	0
	01-FEB-2005 15:40	0

```
                                                01-FEB-2005 16:42          0
                                                01-FEB-2005 16:50          0
```

4.4. the total number of requests for data (over the same period) is:

```
operation class           Total # of requests
------------------------- -------------------
consistent gets                       3878820
db block gets                          218594
```

If the value in 4.3 / the value in 4.4 > 1%, consider increasing number of RBS

Also check the number OLTP users, and number of concurrent transactions.

```
Username.OSUser  SID,Serial# PID       Machine    Active Logon Time         Consis Gets Block Gets Phys Reads Hit ratio
---------------- ----------- --------- ---------- ------ ------------------ ----------- ---------- ---------- ---------
SYS.GChapman     152,101          1996 MK\C-GCHAP No     01-FEB-05 16:04:29        4709      10021       3459     76.52
SYS.GChapman     160,168          3556 MK\C-GCHAP No     01-FEB-05 15:06:49      219785         10      33922     84.57
SYS.GChapman     143,38           4092 MK\C-GCHAP No     01-FEB-05 15:18:48       94857       1046      14290     85.10
SYS.GChapman     142,16           2420 MK\C-GCHAP Yes    01-FEB-05 15:26:32     2497320       4293     209553     91.62
SYS.GChapman     153,328          3096 MK\C-GCHAP No     01-FEB-05 15:26:14       10460      99542       6443     94.14
SYS.GChapman     161,108          3448 MK\C-GCHAP No     01-FEB-05 15:15:41       37414        194       1513     95.98
```

5. Redo log files

Check REDO log files status (look out for any STALE/INVALID/DELETED status)

```
Group# Status  Member
------ ------- ------------------------------------------------
     3         D:\ORACLE\ORADATA\GEN2\REDO03.LOG
     2         D:\ORACLE\ORADATA\GEN2\REDO02.LOG
     1         D:\ORACLE\ORADATA\GEN2\REDO01.LOG
     4         D:\ORACLE\ORADATA\GEN2\REDO04.LOG
```

Info and Average interval for the last 2 days log sequence numbers

```
When                 Interval (mins)         From the           To the
                    between switches         sequence         sequence
-------------------- ---------------- ---------------- ----------------
31-JAN-05 09:54:46              13.42             1113             1114
31-JAN-05 10:08:11               1.00             1114             1115
31-JAN-05 10:09:11              23.23             1115             1116
31-JAN-05 10:32:25              28.78             1116             1117
31-JAN-05 11:01:12              93.42             1117             1118
31-JAN-05 12:34:37              66.02             1118             1119
31-JAN-05 13:40:38              18.83             1119             1120
31-JAN-05 13:59:28               4.52             1120             1121
31-JAN-05 14:03:59               6.48             1121             1122
31-JAN-05 14:10:28               2.73             1122             1123
31-JAN-05 14:13:12               4.67             1123             1124
31-JAN-05 14:17:52               7.32             1124             1125
31-JAN-05 14:25:11              20.35             1125             1126
31-JAN-05 14:45:32                .58             1126             1127
31-JAN-05 14:46:07               7.32             1127             1128
31-JAN-05 14:53:26               7.63             1128             1129
31-JAN-05 15:01:04               8.13             1129             1130
31-JAN-05 15:09:12               4.47             1130             1131
31-JAN-05 15:13:40               7.05             1131             1132
31-JAN-05 15:20:43               8.33             1132             1133
31-JAN-05 15:29:03               1.42             1133             1134
31-JAN-05 15:30:28               5.70             1134             1135
31-JAN-05 15:36:10               3.75             1135             1136
31-JAN-05 15:39:55              16.88             1136             1137
31-JAN-05 15:56:48               3.67             1137             1138
31-JAN-05 16:00:28              11.63             1138             1139
31-JAN-05 16:12:06              29.88             1139             1140
31-JAN-05 16:41:59              52.92             1140             1141
31-JAN-05 17:34:54           1,137.60             1141             1142
01-FEB-05 12:32:30             173.80             1142             1143
01-FEB-05 15:26:18              57.87             1143             1144
```

5.1 Check redo log buffer size.

The REDO BUFFER ALLOCATION RETRIES should be near 0. If it increments consistently, you may need to increase LOG_BUFFER. Note the other factors: checkpointing or archiving.

```
redo buffer allocation retries
------------------------------
                            11
```

```
Activity                                            Log time          Value
--------------------------------------------------- ----------------- ----------
redo buffer allocation retries (Any big increase?)  01-FEB-2005 15:26          6
                                                    01-FEB-2005 15:40          6
                                                    01-FEB-2005 16:42          8
                                                    01-FEB-2005 16:50         11
```

```
 The space request ratio, REDO LOG SPACE REQUESTS / REDO ENTRIES, should
 be < 1:5000 (or 0.02%), otherwise increase the size of the log buffer
 until the ratio stops increasing.

redo log space request ratio %
------------------------------
                          .008


Activity                                               Log time          Value
------------------------------------------------------ ----------------- ----------
redo log space request ratio (should be < 0.02%)       01-FEB-2005 15:26        .00
                                                       01-FEB-2005 15:40        .00
                                                       01-FEB-2005 16:42        .00
                                                       01-FEB-2005 16:50        .00


 5.2  Redo log file wait (contention for I/O) diagnosis:

 Redo log files should be on separate and fast devices because LGWR almost writes non-stop.

 NOTE: The "time_waited" column will be 0 if "timed_statistics" param is FALSE.

# of waits by this operation total wait time
--------------------------- ---------------
                       5955            1335


Activity                                               Log time          Value
------------------------------------------------------ ----------------- ----------
redo log parallel write waits (Any big increase?)      01-FEB-2005 15:26       4278
                                                       01-FEB-2005 15:40       4488
                                                       01-FEB-2005 16:42       5809
                                                       01-FEB-2005 16:50       5955


Activity                                               Log time          Value
------------------------------------------------------ ----------------- ----------
redo log parallel write wait time (Any big increase? 01-FEB-2005 15:26        875
                                                       01-FEB-2005 15:40        901
                                                       01-FEB-2005 16:42       1288
                                                       01-FEB-2005 16:50       1335


 5.3  Redo latch contention

 misses / gets should be < 1% for REDO ALLOCATION latch, and
 immediate_misses / immediate_gets should be < 1% for REDO COPY latch

Name                     Gets    Misses    Sleeps Miss Ratio %
------------------ --------- --------- --------- ------------
redo allocation        23380         1         1          .00


Activity                                               Log time          Value
------------------------------------------------------ ----------------- ----------
redo allocation miss ratio (should be < 1%)            01-FEB-2005 15:26        .00
                                                       01-FEB-2005 15:40        .00
                                                       01-FEB-2005 16:42        .00
                                                       01-FEB-2005 16:50        .00

Name                  Sleeps Immediate Gets Immediate Misses Immediate Miss Ratio %
------------------ --------- -------------- ---------------- ----------------------
redo copy                 0         107064              143                    .13


Activity                                               Log time          Value
------------------------------------------------------ ----------------- ----------
redo copy miss ratio (should be < 1%)                  01-FEB-2005 15:26        .07
                                                       01-FEB-2005 15:40        .09
                                                       01-FEB-2005 16:42        .13
                                                       01-FEB-2005 16:50        .13

 5.4  Redo size used by all transactions


      SID Username        OS User         Logon Time               Redo Size bytes
--------- --------------- --------------- ------------------------ ---------------
      153 SYS             GChapman        01-FEB-05 15:26:14              15943960
      165                 SYSTEM          01-FEB-05 08:02:23               1616040
      152 SYS             GChapman        01-FEB-05 16:04:29               1514036
      158                 SYSTEM          01-FEB-05 08:02:52                582812
      168                 SYSTEM          01-FEB-05 08:02:23                438792
      143 SYS             GChapman        01-FEB-05 15:18:48                286264
      142 SYS             GChapman        01-FEB-05 15:26:32                233476
      161 SYS             GChapman        01-FEB-05 15:15:41                 28516
      163                 SYSTEM          01-FEB-05 08:02:23                  1176
      170                 SYSTEM          01-FEB-05 08:02:23                     0
      146                                 01-FEB-05 08:03:00                     0
      150                                 01-FEB-05 16:00:04                     0
      159                 SYSTEM          01-FEB-05 08:02:45                     0
      141                 SYSTEM          01-FEB-05 16:42:08                     0
      160 SYS             GChapman        01-FEB-05 15:06:49                     0
      162                                 01-FEB-05 16:50:35                     0
      148                                 01-FEB-05 08:03:42                     0
      166                 SYSTEM          01-FEB-05 08:02:23                     0
      169                 SYSTEM          01-FEB-05 08:02:23                     0
      167                 SYSTEM          01-FEB-05 08:02:23                     0
      164                 SYSTEM          01-FEB-05 08:02:23                     0
      156                 SYSTEM          01-FEB-05 08:02:52                     0
                                                                   ---------------
sum                                                                       20645072
```

```
 6. Check the sorts
 ==================

 6.1. Sort-related parameters

Name                                 Value
---------------------------------- ---------------
sort_area_size                           65536
sort_area_retained_size                      0

 disk ratio should be <= 10% and disk sort = 0 (or as near to 0 as possible)

Memory Sort Disk Sort Disk Sort Ratio
------------ --------- ---------------
       54977         0             .00


Activity                                             Log time          Value
---------------------------------------------------- ----------------- ----------
disk/memory sort ratio (should be < 10%)             01-FEB-2005 15:26        .00
                                                     01-FEB-2005 15:40        .00
                                                     01-FEB-2005 16:42        .00
                                                     01-FEB-2005 16:50        .00


 6.2. Monitoring the TEMP datafiles/segments

 sort_area_size is usually set to be == sort_area_retained_size

 Also, the initial_extent and next_extent of the temporary tablespace should be
 multiples of sort_area_size.

Tablespace Name                Initial Extent Next Extent Min Extents Max Extents Pct Increase Status    Contents
------------------------------ -------------- ----------- ----------- ----------- ------------ --------- ---------
TEMP                                  1048576     1048576           1                       0 ONLINE    TEMPORARY

 The total size written to the TEMP datafiles since last DB bounce


DataFile name                               Phy. Size
                                           written Mb
----------------------------------------- ------------

 The space usage of any TEMP segments


Tablespace Segment File Segment Block Extent Size Total Extents Total Blocks Total Size Mb
---------- ------------ ------------- ----------- ------------- ------------ -------------
TEMP                  0             0         128           120        15360             0


 7. Contention in general
 ========================


 The following query generates Wait Statistics.

 This will show wait stats for certain kernal instances - this may show
 the need for additional rollback segements, wait lists, and or
 database buffers.

 NOTE: The "time" column below will be 0 if "timed_statistics" param is FALSE.


 "      Wait Statistics for the Instance"

operation class    # of waits by this operation total wait time
------------------ ---------------------------- ---------------
data block                                  124             236
segment header                                4               2
undo header                                   2               0


 Look at the statistics generated above (if any). They will
 tell the DBA/User where there is contention in the system.
 There will usually be contention in any system - but if the
 ratio of waits for a particular operation starts to rise you may
 need to add additional resource(s), such as more database buffers,
 log buffers or rollback segments.


 8. Check I/O Spread among the DB files (disks)
 ==============================================
 Look for any significant imbalance across disk drives in terms of I/O
 Also examine the Blocks-per-Read RATIO for heavily accessed
 TABLESPACES... If this value is significantly above 1 then you may
 have full tablescans occuring (with multi-block I/O)


                                                          Physical       Physical   Blocks     Physical        Read
Write    Physical Blks
DataFile name                                            Blks Read          Reads Per Read    Blks Wrtn  Time x1000  Time
x1000     Reads+Writes
-------------------------------------------------- -------------- -------------- -------- ------------ ----------- ---------
-- ----------------
1 D:\ORACLE\ORADATA\GEN2\SYSTEM01.DBF                     244,771         50,139        5        1,639          33
2          246,410
3 D:\ORACLE\ORADATA\GEN2\SYSAUX01.DBF                     102,588        102,291        1        4,342          58
5          106,930
4 D:\ORACLE\ORADATA\GEN2\USERS01.DBF                       19,374         19,346        1          285          14
0           19,659
```

```
9 D:\ORACLE\ORADATA\GEN2\MAGTAB01.DBF                       9,564          9,564        1            4           4
0            9,568
5 D:\ORACLE\ORADATA\GEN2\WIXBIL01.DBF                       3,368          3,368        1            4           3
0            3,372
2 D:\ORACLE\ORADATA\GEN2\UNDOTBS01.DBF                        271            271        1        2,610           1
13            2,881
6 D:\ORACLE\ORADATA\GEN2\PMS_DATA01.ORA                     1,376          1,376        1            4           1
0            1,380
12 D:\ORACLE\ORADATA\GEN2\MAG_SIDX1.ORA                     1,112          1,112        1            4           1
0            1,116
7 D:\ORACLE\ORADATA\GEN2\PMS_IDX.ORA                          980            980        1            4           1
0              984
11 D:\ORACLE\ORADATA\GEN2\CAPTURE_IDX1.ORA                    408            408        1            4           0
0              412
10 D:\ORACLE\ORADATA\GEN2\CAPTURE.ORA                         232            232        1            4           0
0              236
```

8.1. Show I/O distribution across disk drives

```
Datafile name                       Total reads+writes Size of all DB files (Mb)
----------------------------------- ------------------ --------------------------
D:\ORACLE\ORADATA\GEN2                          195491                       2160
```

8.2. Check the "last backup" timestamp of all the datafiles

```
Last Backup date       # of datafiles
---------------------- --------------
NEVER                              11
```

A list of any datafiles whose last backup is older than 2 days ago.

```
File name                                 Ts Name      Date                Status     Change#
----------------------------------------- ------------ ------------------- ---------- --------
D:\ORACLE\ORADATA\GEN2\SYSTEM01.DBF       SYSTEM       NEVER               NOT ACTIVE        0
D:\ORACLE\ORADATA\GEN2\UNDOTBS01.DBF      UNDOTBS1     NEVER               NOT ACTIVE        0
D:\ORACLE\ORADATA\GEN2\SYSAUX01.DBF       SYSAUX       NEVER               NOT ACTIVE        0
D:\ORACLE\ORADATA\GEN2\USERS01.DBF        USERS        NEVER               NOT ACTIVE        0
D:\ORACLE\ORADATA\GEN2\WIXBIL01.DBF       WIXBIL       NEVER               NOT ACTIVE        0
D:\ORACLE\ORADATA\GEN2\PMS_DATA01.ORA     PMS_DATA     NEVER               NOT ACTIVE        0
D:\ORACLE\ORADATA\GEN2\PMS_IDX.ORA        PMS_IDX      NEVER               NOT ACTIVE        0
D:\ORACLE\ORADATA\GEN2\MAGTAB01.DBF       MAGELLAN     NEVER               NOT ACTIVE        0
D:\ORACLE\ORADATA\GEN2\CAPTURE.ORA        CAPTURE      NEVER               NOT ACTIVE        0
D:\ORACLE\ORADATA\GEN2\CAPTURE_IDX1.ORA   CAPTURE_IDX  NEVER               NOT ACTIVE        0
D:\ORACLE\ORADATA\GEN2\MAG_SIDX1.ORA      MAG_SIDX     NEVER               NOT ACTIVE        0
```

9. Any space-critical segments?
==========

Show segments whose next 5 extents can't fit in the SUM of all
the free space in a tablespace.

```
**************************************************************************************************
Owner      Tablespace Name        Object Name                    Obj    Next Extent  Space Left Biggest Chunk            Space
Used
                                                                 Type       Size Mb  for # EXTs            Mb   HWM Mb
Mb    Extents
---------- ---------------------- ------------------------------ ------ ---------- ---------- ------------- --------- -----
----- ----------
**************************************************************************************************
```

9.1 Any space-critical segments? Archive tables/indexes only

Show segments whose next 5 extents can't fit in the SUM of all
the free space in a tablespace.

```
**************************************************************************************************
Owner      Tablespace Name        Object Name                    Obj    Next Extent  Space Left Biggest Chunk            Space
Used
                                                                 Type       Size Mb  for # EXTs            Mb   HWM Mb
Mb    Extents
---------- ---------------------- ------------------------------ ------ ---------- ---------- ------------- --------- -----
----- ----------
**************************************************************************************************
```

9.2 Any new objects been created.

```
**************************************************************************************************
Object                         Owner      Type         Status  Created
------------------------------ ---------- ------------ ------- -------------------
EMAIL                          GSC        PACKAGE BODY VALID   2005-02-01:15:44:04
LOGDIR                         SYS        DIRECTORY    VALID   2005-02-01:15:49:55
EMAIL                          GSC        PACKAGE      VALID   2005-02-01:12:21:50
```

9.3 Check whether the number of datafiles approaches its ceiling?
NOTE that 'Max # of datafiles' below is the value of DB_FILES in your init.ora,
it is <= the MAXDATAFILES which can be found out in a Controlfile Trace.

```
**************************************************************************************************
Current # of datafiles   Max # of datafiles   Remaining file slots
                         in this instance     in this instance
------------------------ -------------------- --------------------
11                       200                                  189
```

10. License Usage Details (current and maximum # of sessions)

```
Sessions Max Sessions_Warning Sessions Current Sessions Highwater
------------ ---------------- ---------------- ------------------
           0                0               10                 10
**************************************************************************************************
```

```
 ********* End of the check

Db Check Ends: (DD-MON-RR HH24:MI:SS)
-------------------------------------
01-FEB-05 16:50:52
```

C.2 Sample ADDM Output

```
          DETAILED ADDM REPORT FOR TASK 'TASK_3068' WITH ID 3068
          ------------------------------------------------------

              Analysis Period: from 29-JAN-2005 00:00 to 31-JAN-2005 11:00
         Database ID/Instance: 2654633518/1
      Database/Instance Names: SHQCTMS/shqctms
                    Host Name: SHQCTMSDB
            Database Version: 10.1.0.3.0
              Snapshot Range: from 3010 to 3069
                Database Time: 2062 seconds
        Average Database Load: 0 active sessions

~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~


FINDING 1: 38% impact (789 seconds)
-----------------------------------
Time spent on the CPU by the instance was responsible for a substantial part
of database time.

   RECOMMENDATION 1: SQL Tuning, 17% benefit (357 seconds)
      ACTION: Tune the PL/SQL block with SQL_ID "2b064ybzkwf1y". Refer to the
         "Tuning PL/SQL Applications" chapter of Oracle's "PL/SQL User's Guide
         and Reference"
         RELEVANT OBJECT: SQL statement with SQL_ID 2b064ybzkwf1y
         BEGIN EMD_NOTIFICATION.QUEUE_READY(:1, :2, :3); END;

   RECOMMENDATION 2: SQL Tuning, 6.9% benefit (142 seconds)
      ACTION: Run SQL Tuning Advisor on the SQL statement with SQL_ID
         "8hk7xvhua40va".
         RELEVANT OBJECT: SQL statement with SQL_ID 8hk7xvhua40va and
         PLAN_HASH 1
         INSERT INTO MGMT_METRICS_RAW(COLLECTION_TIMESTAMP, KEY_VALUE,
         METRIC_GUID, STRING_VALUE, TARGET_GUID, VALUE) VALUES ( :1, NVL(:2, '
         '), :3, :4, :5, :6)

   RECOMMENDATION 3: SQL Tuning, 5.2% benefit (107 seconds)
      ACTION: Tune the PL/SQL block with SQL_ID "cb75rw3w1tt0s". Refer to the
         "Tuning PL/SQL Applications" chapter of Oracle's "PL/SQL User's Guide
         and Reference"
         RELEVANT OBJECT: SQL statement with SQL_ID cb75rw3w1tt0s
         begin MGMT_JOB_ENGINE.get_scheduled_steps(:1, :2, :3, :4); end;

   RECOMMENDATION 4: SQL Tuning, 4.7% benefit (97 seconds)
      ACTION: Tune the PL/SQL block with SQL_ID "010ycx2rw8tnc". Refer to the
         "Tuning PL/SQL Applications" chapter of Oracle's "PL/SQL User's Guide
         and Reference"
         RELEVANT OBJECT: SQL statement with SQL_ID 010ycx2rw8tnc
         begin
         dbms_stats.gather_schema_stats(ownname=> 'ctms' , estimate_percent =>
         NULL, cascade=> TRUE);
         end;

   RECOMMENDATION 5: SQL Tuning, 4.1% benefit (85 seconds)
      ACTION: Tune the PL/SQL block with SQL_ID "6q766vsk5290x". Refer to the
         "Tuning PL/SQL Applications" chapter of Oracle's "PL/SQL User's Guide
         and Reference"
         RELEVANT OBJECT: SQL statement with SQL_ID 6q766vsk5290x
         BEGIN dbms_stats_internal.flush_cache_stats; END;
```

```
FINDING 2: 29% impact (594 seconds)
-----------------------------------
Wait class "Other" was consuming significant database time.

   NO RECOMMENDATIONS AVAILABLE

   ADDITIONAL INFORMATION: Database latches in the "Other" wait class were not
      consuming significant database time.

FINDING 3: 29% impact (593 seconds)
-----------------------------------
Wait event "class slave wait" in wait class "Other" was consuming significant
database time.

   RECOMMENDATION 1: Application Analysis, 29% benefit (593 seconds)
      ACTION: Investigate the cause for high "class slave wait" waits. Refer
         to Oracle's "Database Reference" for the description of this wait
         event. Use given SQL for further investigation.
      RATIONALE: The SQL statement with SQL_ID "NULL-SQLID" was found waiting
         for "class slave wait" wait event.
         RELEVANT OBJECT: SQL statement with SQL_ID NULL-SQLID


   RECOMMENDATION 2: Application Analysis, 29% benefit (593 seconds)
      ACTION: Investigate the cause for high "class slave wait" waits in
         Service "SYS$BACKGROUND".

FINDING 4: 25% impact (513 seconds)
-----------------------------------
SQL statements consuming significant database time were found.

   RECOMMENDATION 1: SQL Tuning, 8% benefit (164 seconds)
      ACTION: Tune the PL/SQL block with SQL_ID "010ycx2rw8tnc". Refer to the
         "Tuning PL/SQL Applications" chapter of Oracle's "PL/SQL User's Guide
         and Reference"
         RELEVANT OBJECT: SQL statement with SQL_ID 010ycx2rw8tnc
         begin
         dbms_stats.gather_schema_stats(ownname=> 'ctms' , estimate_percent =>
         NULL, cascade=> TRUE);
         end;

   RECOMMENDATION 2: SQL Tuning, 6.3% benefit (131 seconds)
      ACTION: Run SQL Tuning Advisor on the SQL statement with SQL_ID
         "8hk7xvhua40va".
         RELEVANT OBJECT: SQL statement with SQL_ID 8hk7xvhua40va and
         PLAN_HASH 1
         INSERT INTO MGMT_METRICS_RAW(COLLECTION_TIMESTAMP, KEY_VALUE,
         METRIC_GUID, STRING_VALUE, TARGET_GUID, VALUE) VALUES ( :1, NVL(:2, '
         '), :3, :4, :5, :6)

   RECOMMENDATION 3: SQL Tuning, 4.2% benefit (86 seconds)
      ACTION: Tune the PL/SQL block with SQL_ID "6q766vsk5290x". Refer to the
         "Tuning PL/SQL Applications" chapter of Oracle's "PL/SQL User's Guide
         and Reference"
         RELEVANT OBJECT: SQL statement with SQL_ID 6q766vsk5290x
         BEGIN dbms_stats_internal.flush_cache_stats; END;

   RECOMMENDATION 4: SQL Tuning, 3.3% benefit (69 seconds)
      ACTION: Tune the PL/SQL block with SQL_ID "cb75rw3w1tt0s". Refer to the
         "Tuning PL/SQL Applications" chapter of Oracle's "PL/SQL User's Guide
         and Reference"
         RELEVANT OBJECT: SQL statement with SQL_ID cb75rw3w1tt0s
         begin MGMT_JOB_ENGINE.get_scheduled_steps(:1, :2, :3, :4); end;

   RECOMMENDATION 5: SQL Tuning, 3.1% benefit (63 seconds)
      ACTION: Tune the PL/SQL block with SQL_ID "6gvch1xu9ca3g". Refer to the
         "Tuning PL/SQL Applications" chapter of Oracle's "PL/SQL User's Guide
         and Reference"
```

```
         RELEVANT OBJECT: SQL statement with SQL_ID 6gvch1xu9ca3g
         DECLARE job BINARY_INTEGER := :job; next_date DATE := :mydate;
         broken BOOLEAN := FALSE; BEGIN
         EMD_MAINTENANCE.EXECUTE_EM_DBMS_JOB_PROCS(); :mydate := next_date; IF
         broken THEN :b := 1; ELSE :b := 0; END IF; END;

FINDING 5: 18% impact (375 seconds)
-----------------------------------
PL/SQL execution consumed significant database time.

   RECOMMENDATION 1: SQL Tuning, 18% benefit (375 seconds)
      ACTION: Tune the PL/SQL block with SQL_ID "2b064ybzkwf1y". Refer to the
         "Tuning PL/SQL Applications" chapter of Oracle's "PL/SQL User's Guide
         and Reference"
         RELEVANT OBJECT: SQL statement with SQL_ID 2b064ybzkwf1y
         BEGIN EMD_NOTIFICATION.QUEUE_READY(:1, :2, :3); END;

FINDING 6: 5.3% impact (109 seconds)
------------------------------------
SQL statements were not shared due to the usage of literals. This resulted in
additional hard parses which were consuming significant database time.

   RECOMMENDATION 1: Application Analysis, 5.3% benefit (109 seconds)
      ACTION: Investigate application logic for possible use of bind variables
         instead of literals. Alternatively, you may set the parameter
         "cursor_sharing" to "force".
      RATIONALE: SQL statements with PLAN_HASH_VALUE 687186664 were found to
         be using literals. Look in V$SQL for examples of such SQL statements.
      RATIONALE: SQL statements with PLAN_HASH_VALUE 216435276 were found to
         be using literals. Look in V$SQL for examples of such SQL statements.
      RATIONALE: SQL statements with PLAN_HASH_VALUE 1586278049 were found to
         be using literals. Look in V$SQL for examples of such SQL statements.
      RATIONALE: SQL statements with PLAN_HASH_VALUE 3887861132 were found to
         be using literals. Look in V$SQL for examples of such SQL statements.

   SYMPTOMS THAT LED TO THE FINDING:
      Hard parsing of SQL statements was consuming significant database time.
      (9.2% impact [189 seconds])

FINDING 7: 3.8% impact (79 seconds)
-----------------------------------
Cursors were getting invalidated due to DDL operations. This resulted in
additional hard parses which were consuming significant database time.

   RECOMMENDATION 1: Application Analysis, 3.8% benefit (79 seconds)
      ACTION: Investigate appropriateness of DDL operations.

   SYMPTOMS THAT LED TO THE FINDING:
      Hard parsing of SQL statements was consuming significant database time.
      (9.2% impact [189 seconds])

FINDING 8: 3.7% impact (75 seconds)
-----------------------------------
Individual SQL statements responsible for significant user I/O wait were
found.

   RECOMMENDATION 1: SQL Tuning, 3.6% benefit (75 seconds)
      ACTION: Tune the PL/SQL block with SQL_ID "010ycx2rw8tnc". Refer to the
         "Tuning PL/SQL Applications" chapter of Oracle's "PL/SQL User's Guide
         and Reference"
         RELEVANT OBJECT: SQL statement with SQL_ID 010ycx2rw8tnc
         begin
         dbms_stats.gather_schema_stats(ownname=> 'ctms' , estimate_percent =>
         NULL, cascade=> TRUE);
         end;

   RECOMMENDATION 2: SQL Tuning, 0.01% benefit (0 seconds)
      ACTION: Tune the PL/SQL block with SQL_ID "6gvch1xu9ca3g". Refer to the
```

```
          "Tuning PL/SQL Applications" chapter of Oracle's "PL/SQL User's Guide
          and Reference"
          RELEVANT OBJECT: SQL statement with SQL_ID 6gvch1xu9ca3g
          DECLARE job BINARY_INTEGER := :job; next_date DATE := :mydate;
          broken BOOLEAN := FALSE; BEGIN
          EMD_MAINTENANCE.EXECUTE_EM_DBMS_JOB_PROCS(); :mydate := next_date; IF
          broken THEN :b := 1; ELSE :b := 0; END IF; END;

   RECOMMENDATION 3: SQL Tuning, 0% benefit (0 seconds)
      ACTION: Tune the PL/SQL block with SQL_ID "6q766vsk5290x". Refer to the
         "Tuning PL/SQL Applications" chapter of Oracle's "PL/SQL User's Guide
         and Reference"
         RELEVANT OBJECT: SQL statement with SQL_ID 6q766vsk5290x
         BEGIN dbms_stats_internal.flush_cache_stats; END;

   RECOMMENDATION 4: SQL Tuning, 0% benefit (0 seconds)
      ACTION: Run SQL Tuning Advisor on the SQL statement with SQL_ID
         "8hk7xvhua40va".
         RELEVANT OBJECT: SQL statement with SQL_ID 8hk7xvhua40va and
         PLAN_HASH 1
         INSERT INTO MGMT_METRICS_RAW(COLLECTION_TIMESTAMP, KEY_VALUE,
         METRIC_GUID, STRING_VALUE, TARGET_GUID, VALUE) VALUES ( :1, NVL(:2, '
         '), :3, :4, :5, :6)

   RECOMMENDATION 5: SQL Tuning, 0% benefit (0 seconds)
      ACTION: Tune the PL/SQL block with SQL_ID "cb75rw3w1tt0s". Refer to the
         "Tuning PL/SQL Applications" chapter of Oracle's "PL/SQL User's Guide
         and Reference"
         RELEVANT OBJECT: SQL statement with SQL_ID cb75rw3w1tt0s
         begin MGMT_JOB_ENGINE.get_scheduled_steps(:1, :2, :3, :4); end;

   SYMPTOMS THAT LED TO THE FINDING:
      Wait class "User I/O" was consuming significant database time. (6%
      impact [124 seconds])

FINDING 9: 2.5% impact (51 seconds)
-----------------------------------
Soft parsing of SQL statements was consuming significant database time.

   RECOMMENDATION 1: Application Analysis, 2.5% benefit (51 seconds)
      ACTION: Investigate application logic to keep open the frequently used
         cursors. Note that cursors are closed by both cursor close calls and
         session disconnects.

   RECOMMENDATION 2: DB Configuration, 2.5% benefit (51 seconds)
      ACTION: Consider increasing the maximum number of open cursors a session
         can have by increasing the value of parameter "open_cursors".
      ACTION: Consider increasing the session cursor cache size by increasing
         the value of parameter "session_cached_cursors".

FINDING 10: 2% impact (41 seconds)
----------------------------------
The PGA was inadequately sized, causing additional I/O to temporary
tablespaces to consume significant database time.

   RECOMMENDATION 1: DB Configuration, 1.6% benefit (34 seconds)
      ACTION: Increase the size of the PGA by setting the value of parameter
         "pga_aggregate_target" to 240 M.

   SYMPTOMS THAT LED TO THE FINDING:
      Wait class "User I/O" was consuming significant database time. (6%
      impact [124 seconds])


~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~~

          ADDITIONAL INFORMATION
          ----------------------
```

```
Wait class "Administrative" was not consuming significant database time.
Wait class "Application" was not consuming significant database time.
Wait class "Cluster" was not consuming significant database time.
Wait class "Commit" was not consuming significant database time.
Wait class "Concurrency" was not consuming significant database time.
Wait class "Configuration" was not consuming significant database time.
Wait class "Network" was not consuming significant database time.
Wait class "Scheduler" was not consuming significant database time.

The analysis of I/O performance is based on the default assumption that the
average read time for one database block is 10000 micro-seconds.

An explanation of the terminology used in this report is available when you
run the report with the 'ALL' level of detail.
```

D. Installation script

```
rem
rem Script to create the stats pkg owner.
rem
rem Version 1.2   8-Feb-05  GSC
rem
prompt 'This script needs to run as the sys user to ensure'
prompt' that the corect permissions are given.'
prompt
prompt 'Please enter stats package owner:'
accept statsown
prompt ‘Please enter stats package owner default tablespace’
accept deftblspce

prompt 'Dropping stats package owner if they exist.'
drop user &statsown cascade;
create user &statsown identified by abcd default tablespace &deftblspce temporary tablespace
temp;
alter user &statsown quota unlimited on &deftblspce;
GRANT CONNECT,RESOURCE TO &statsown;
GRANT CREATE TABLE TO &statsown;
grant create procedure to &statsown;

REM Following required for access to underlying SYS views
grant select on dba_tab_columns to &statsown;
grant select on dba_tables to &statsown;
grant select on dba_users to &statsown;
grant select on dba_objects to &statsown;

grant select on ts$ to &statsown;
grant select on dba_tables to &statsown;
grant select on dba_free_space to &statsown;
grant select on dba_extents to &statsown;
grant select on dba_segments to &statsown;
grant select on v_$filestat to &statsown;
grant select on v_$dbfile to &statsown;
grant select on v_$loghist to &statsown;
grant select on v_$session to &statsown;
grant select on v_$sess_io to &statsown;
grant select on v_$process to &statsown;
grant select on v_$db_object_cache to &statsown;
grant select on v_$parameter to &statsown;
grant select on v_$license to &statsown;
grant select on dba_data_files to &statsown;
grant select on dba_tablespaces to &statsown;
grant select on v_$backup to &statsown;
grant select on v_$waitstat to &statsown;
grant select on v_$sort_segment to &statsown;
grant select on v_$database to &statsown;
grant select on dba_tab_columns to &statsown;
grant select on v_$sysstat to &statsown;
grant select on v_$sesstat to &statsown;
grant select on v_$session to &statsown;
grant select on v_$logfile to &statsown;
grant select on v_$rollstat to &statsown;
grant select on v_$rollname to &statsown;
grant select on v_$latch to &statsown;
grant select on v_$librarycache to &statsown;
grant select on v_$rowcache to &statsown;
grant select on v_$system_event to &statsown;
grant select on dba_hist_snapshot to &statsown;
grant select on v_$instance to &statsown;
```

```
rem Following required for dbms_scheduler
GRANT CREATE job TO &statsown;
GRANT EXECUTE ANY program TO &statsown;
GRANT EXECUTE ANY class TO &statsown;
GRANT manage scheduler TO &statsown;

rem Following required for AWR procedure
grant execute on dbms_workload_manager to &statsown;

rem Following required for ADDM
grant select on dba_advisor_tasks to &statsown;
GRANT advisor TO &statsown;
GRANT EXECUTE ON dbms_advisor TO &statsown;

GRANT CREATE ANY DIRECTORY TO &statsown;
prompt
prompt Now I need TO CREATE a DIRECTORY
prompt 'Please enter the full directory name where logs are to be placed '
ACCEPT logspec

CREATE DIRECTORY logdir AS &logspec;
prompt 'Changing to stats Table owner'
connect &statsown/abcd
prompt 'Creating directory '
CREATE DIRECTORY LOGDIR AS '&logspec';

prompt 'Generating the packages.'
prompt
@log_spec.SQL
@log_body.SQL
@email_spec.SQL
@email_body.SQL
@statgen_spec.SQL
@statgen_body.SQL
rem
rem Now generate the TABLES.
rem
BEGIN
  stats_gen.gen_tables('&deftblspce');
  stats_gen.gen_dba_tab_growth('&deftblspce');
END;
/
@stats_spec.SQL
@stats_body.SQL

prompt
prompt 'Now we must change the audit owners password.'
prompt 'Please enter the required password.'
accept passwd
alter user &statsown identified by &passwd;
prompt
prompt ' Now check that everything is fine and you can run '
prompt ' the statistics report'
prompt
rem exit;
```

E. Outstanding Activities

E.1 Oracle 10g dependencies

Despite earlier statements about running upon all versions of Oracle, the recent code changes have introduced reliance upon Oracle 10g. It is true that the basic STATS procedures remain valid for pre Oracle 10g databases. The use of features such as AWR, ADDM and DBMS_SCHEDULER introduce the Oracle 10g dependencies. These can/could be changed should need arise. The AWR reports could be substituted by Statspack. ADDM unfortunately has no substitution, and DBMS_SCHEDULER could be replaced by the use of DBMS_JOB.

A further dependencies lies in a few of the procedures with their use of NDS which was introduced with Oracle 8.1. Hence this will generate difficulties with Oracle 7 or Oracle 8. The use of DBMS_parse could be substituted if required, but since the use of NDS is related to the Oracle 01g features such as ADDM, AWR etc, there is probably not much to be gained by using dbms_parse.

E.2 Email considerations

Oracle 8.1.6 introduced UTL_TCP into the database, built on top of Java sockets. The EMAIL package uses UTL_SMTP, which in turn calls UTL_TCP procedures to send email. Oracle 10g introduced a package UTL_MAIL that could be used, but would introduce an additional Oracle 10g dependency. Consequently it has not been used.

It is also possible to load JavaMail into the database to provide a 'richer' email environment. This work has not been done for the current implementation, for which it would be overkill. Preliminary investigations seem to indicate that in Oracle 10g there is a 'sys' implementation of javamail, but no work has been performed to make user procedures call to the underlying JAVA code.

E.3 Field values overflowing

For a database that has been running for a long period of time or with a lot of activity, because the statistics collected are cumulative, there are some fields in the output displays with are displayed as a string of asterisks '*******'. This is not a problem per sec, but an indication that the length of the value is such that the display can not show the value. There is no real solution to this problem other than increasing the display size of the required field within the Stats package.

Some of the 'system' generated AWR reports are also known to show this behaviour.